

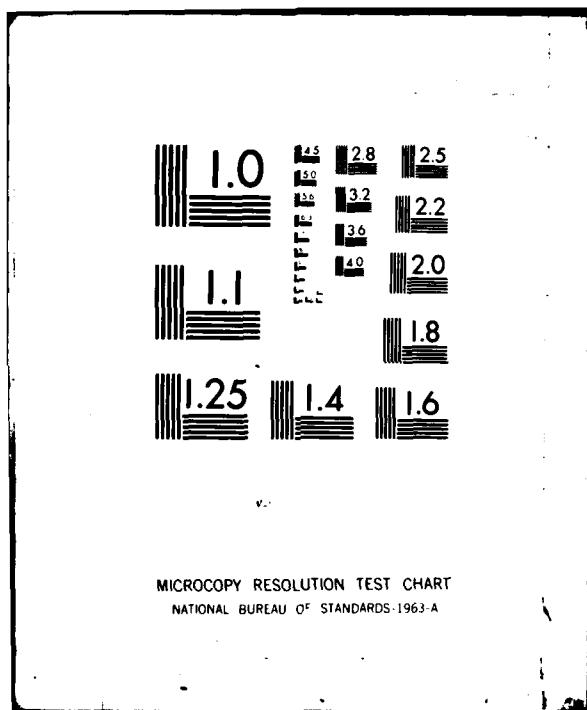
AD-A112 566 RAND CORP SANTA MONICA CA
SCENARIOS FOR EVOLUTION OF AIR TRAFFIC CONTROL. (U)
NOV 81 R WESSON, K SOLOMON, R STEEB
UNCLASSIFIED RAND/R-2698-FAA

F/G 17/7

MDA903-81-C-0211

NL

END
4 82
DTIC



(12)

AD A112566

Scenarios for Evolution of Air Traffic Control

**Robert Wesson, Kenneth Solomon,
Randall Steeb, Perry Thorndyke,
and Keith Wescourt**

DTIC FULL COPY



Rand

The research described in this report was sponsored by the Federal Aviation Administration under Contract No. MDA903-81-C-0211.

Library of Congress Cataloging in Publication Data

Main entry under title:

Scenarios for evolution of air traffic control.

"R-2698-FAA."

I. Air traffic control. I. Wesson, Robert.
II. United States. Federal Aviation Administration.
III. Rand Corporation.
TL725.3.T7S25 629.1,.6'6 81-15808
ISBN 0-8330-0360-7 AACR2

The Rand Publication Series: The Report is the principal publication documenting and transmitting Rand's major research findings and final research results. The Rand Note reports other outputs of sponsored research for general distribution. Publications of The Rand Corporation do not necessarily reflect the opinions or policies of the sponsors of Rand research.

Published by The Rand Corporation

R-2698-FAA

Scenarios for Evolution of Air Traffic Control

**Robert Wesson, Kenneth Solomon,
Randall Steeb, Perry Thorndyke,
and Keith Wescourt**

November 1981

**Prepared for the
Federal Aviation Administration**



Accession For	
NTIS	GRA&I <input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By _____	
Distribution/	
Availability Codes	
Dist	Avail. for Spec.
A	

APPROVED FOR PUBLIC RELEASE: DISTRIBUTION UNLIMITED

PREFACE

Over the past two decades, computers have been playing an increasingly important role in Air Traffic Control (ATC) in the United States. Despite this progress, the process of ATC and particularly the decision-making role remain the responsibility of the air traffic controller. He is still responsible for the second-to-second control of aircraft.

Advances in computer hardware and software technology now promise greater automation of the ATC process and a significantly different role for the controller. This automation brings with it the prospect of greater productivity for controllers and more fuel-efficient flight. Since the mid-1970s, the Federal Aviation Administration (FAA) has been exploring means of achieving these promised benefits. Laboratory simulations have demonstrated that computers can be programmed to generate fuel-efficient, conflict-free flight profiles and the necessary aircraft clearances (i.e., commands) for automatic transmission to pilots.

In 1979, the FAA's confidence in the success of higher levels of automation and the strong support of the ATC system user community accelerated the pursuit of a more automated ATC system. A small team of industry and FAA experts was assembled to develop a concept for Automated En Route ATC (AERA). The results of that effort are documented in *The AERA Concept* (FAA-EM-81-3). At about the same time, a project sponsored by the FAA was undertaken at The Rand Corporation to consider alternative scenarios for evolution to a highly automated ATC system. An interdisciplinary team of Rand computer scientists, engineers, and psychologists concentrated on the relative roles of the controller and the computer and, more specifically, on the preferred interactions between man and machine.

Uncertainties in the human's proposed role under the AERA concept prompted the preliminary design of a variety of alternatives to AERA. This report describes and compares the critical human-factors problems involved in AERA and a particular alternative called Shared Control.¹ The results were generated by applying the somewhat limited existing body of knowledge on human factors in man/computer interactions to future ATC concepts that substantially exceed current experience in terms of task complexity and level of automation. The

¹The other alternative ATC concepts that were constructed are described in the Appendixes to this report.

analyses reported here and the conclusions regarding the advantages and disadvantages of each concept must therefore be presented mainly in qualitative terms.

The FAA is now planning research, development, and experimentation that will carry on this effort. Future work will be directed toward questions that still remain open—work that will help define the ATC system for the year 2000 and will identify appropriate paths for evolution to that system.

SUMMARY

To accommodate the predicted demand for air traffic service in the year 2000, computer technology must augment human control skills. Preliminary laboratory studies have demonstrated that computer programs can track aircraft, predict their future paths, generate conflict-free clearances, and monitor them for compliance—all automatically. This technology could automate most routine ATC tasks and could change the human role in ATC to that of a system manager. How to make the transition to such a system from the present one and exactly what the future specialist's role would be are the issues addressed by this report.

We present three scenarios that delineate a spectrum of transition plans: a Baseline scenario in which the human controller's role is emphasized; an AERA (Automated En Route ATC) scenario in which computers assume the primary control responsibility and perform most ATC functions autonomously; and a Shared Control scenario in which automated, individually invokable modules assist a human specialist who retains the primary responsibility for control.

We compare each scenario's potential for meeting three objectives: increased safety, increased fuel efficiency, and increased controller productivity. Our analytic framework rests on four principles: cost effectiveness, technical conservatism, evolutionary progress, and human involvement.

The Baseline scenario ultimately is uninteresting because its "business as usual" philosophy leads to greatly increased staffing costs to pay for reduced performance. Projected increases in demand for ATC services will increase controller workloads and reduce margins for error. Adding more controllers and reducing sector size may meet this demand temporarily, but increases in intersector coordination requirements, communication channel overload, and human cognitive limitations will tend to reduce overall system safety and performance over the longer term.

The AERA scenario culminates in a very highly automated ATC system by the year 2000. This system would automatically perform most control functions in en route high-altitude and transition sectors. Because an AERA system would operate almost autonomously, with its human "system managers" outside the routine time-critical control loop, it requires virtually perfect software and a complex fail-safe design. If AERA can be realized, its limited domain of applicability and

lengthy development time frame are likely to greatly reduce its potential gains. Much greater technological risks would be incurred in developing the AERA concept than in developing the other concepts addressed here.

The Shared Control scenario offers a compromise between Baseline and AERA by implementing modules similar to those of AERA as controller aids at regular intervals. During the 1980s, for example, digital communications with a tactical communications management software system would enable controllers to store planned clearances for later automatic delivery. Strategic and tactical planning aids, combined with track monitoring aids, would extend their visualization abilities and allow more fuel-efficient clearances. Later, during the 1990s, these functions could be integrated by an executive module. However, unlike the AERA system, this module would perform only fill-in duties for the controller. In the Shared Control scenario, basic separation-assurance responsibility is assigned to the machine (which continuously checks tracks for possible conflicts and intervenes with avoidance instructions if required). The human controller remains firmly in command of his suite of automated tools.

The aiding modules of Shared Control should be applicable to more (and more problematical) ATC domains than the positive-control airspaces of the AERA concept. They should enable future controllers to provide better dissimilar redundancy for the ATC system. Consequently, manning requirements would be limited while the system evolves gradually into a highly automated year-2000 ATC system comparable in capability to AERA, but quite different in its proposed human role.

Except for these role differences and the manner in which individual modules of automation are deployed and integrated, the Shared Control and AERA scenarios differ very little. To combine the two concepts, it would be necessary only to replace AERA's emphasis on automating as much of ATC as possible with Shared Control's emphasis on extending human capabilities through a series of evolutionary automated aids.

CONTENTS

PREFACE	iii
SUMMARY	v
FIGURES AND TABLES	ix
Section	
I. INTRODUCTION	1
II. ALTERNATIVE ATC SCENARIOS	4
Baseline	4
AERA	8
Shared Control	11
Other Scenarios: High-Technology ATC	15
III. GUIDING PRINCIPLES	18
Cost Effectiveness	18
Technical Conservatism	19
Evolutionary Progress	19
Human Involvement	20
IV. HUMAN ROLES IN EACH SCENARIO	25
Baseline	25
AERA	26
Shared Control	35
V. TECHNICAL AND ECONOMIC IMPLICATIONS	39
Baseline	39
AERA	42
Shared Control	46
Summary of the Technological Issues	50
VI. CONCLUSIONS	56
Appendix	
A. SATELLITE-BASED AIR TRAFFIC CONTROL	59
B. ELECTRONIC FLIGHT RULE AIR TRAFFIC CONTROL	62
C. EXPERIMENTS TO REDUCE THE UNCERTAINTIES	66
REFERENCES	71

FIGURES

2.1.	Major AERA Functions	9
2.2.	Shared Control Final System Module Configuration	16
5.1.	Actual/Projected Controller Productivity, 1960 to 2000 ..	41
5.2.	Anticipated Staffing Levels Under Alternative Scenarios.	49

TABLES

2.1.	Synopsis of AERA Scenario Main Events	12
2.2.	Synopsis of Shared Control Scenario Main Events	17
4.1.	Human Roles in AERA	28

I. INTRODUCTION

This nation's air traffic control (ATC) system facilitates the movement of thousands of aircraft every day. It has grown from a few independent radar systems in the early 1950s into a highly sophisticated hierarchical network composed of control towers, approach controls, and en route control facilities. From takeoff to touchdown, many private and all commercial and military flight operations depend on radar-assisted separation and flow management by ATC personnel.

The ATC system of today rarely fails to provide safe and expeditious movement of aircraft. It enables flights to operate in virtually all weather conditions. Using computers, radar, and a cadre of highly skilled controllers, the system assures pilots of adequate separation from one another even when they cannot see beyond the windshield. It manages the limited capacities of our airports, impartially merging small private planes into the same landing patterns that serve jumbo jets.

To function, this system depends primarily on the successful interplay of man and machine. While controllers in the smaller-airport tower cabs may be assisted only by a simple VHF radio, most control functions require the technology of modern electronics—radio, radar, and computers. Darkened rooms full of humming equipment house row upon row of radar-generated, computer-enhanced video displays. ATC computers associate the radar blips with stored flight-plan data, tag each target on the scope with its identity and altitude, and continuously check for conflicting situations in which aircraft may pass too close together or descend too low.

More computer technology is on the way. Microwave landing systems, new controller displays, and new collision-alert systems are among the new electronic tools now under active development by a far-reaching, FAA-sponsored R&D program. As more of these systems become available, controllers and pilots alike will depend more heavily on them. Inevitably, human skills are giving way to automated control systems.

Increased automation will help to achieve the three primary goals of ATC:

- Increased safety.
- Aircraft operation along optimal fuel-efficient profiles with minimal interference.
- Increased productivity of individual controllers.

To achieve these goals, it will be necessary to overcome the limitations in the present ATC system and the problems it faces. Human eyes do not see through fog; human minds sometimes make dangerous mistakes which may be caught by an automated backup system. More crowded skies mean more procedural constraints on aircraft profiles imposed by human controllers. They also mean more competition for the same limited resources and less margin for error. The 62 percent increase in sheer numbers of aircraft operations predicted by the FAA for 1992 means more sectors, more controllers, more coordination, and more dollars, unless productivity can be increased proportionately.

One way to meet this increased demand is to develop a very highly automated ATC system. The prospect of almost total automation is no longer only science fiction. Computers are powerful and fast enough to project aircraft flight paths far into the future, to automatically correct them when they conflict with the anticipated flight profiles of nearby aircraft, and to digitally transmit the revised clearances up to the aircraft. Machines can continuously compute and update delay predictions, so that aircraft can be slowed at fuel-efficient higher altitudes when airports are operating at peak capacities. FAA-sponsored laboratory research is in fact already laying the foundation for a future, very highly automated ATC system.

The critical question in designing the ATC system of the future is not really what *can* be done but what *should* be done. Exactly how much and what kind of automation should assist or replace the human controller? Should we strive for a system in which the machine has the primary responsibility of control and human expertise is used in a secondary, backup fashion? Or should men, in spite of their intrinsic limitations, retain primary control responsibility and utilize machine aids to extend their abilities? Just what is man's optimal role in a highly automated ATC system?

Once a future system is designed, another set of troublesome questions concerns how to implement it. What development and deployment hurdles stand in its way? What are the best evolutionary pathways from ATC circa 1980 to ATC circa 2000? What are the options and what costs and benefits must be carefully balanced before choosing among them?

This report proposes a few possible alternative pathways for ATC evolution during the next two decades or so. We examine, compare, and criticize these alternatives, using various metrics. We discuss their various advantages and disadvantages. We must emphasize that our examination, comparison, and criticism do not take the form of a traditional analysis. Quantitative "hard" data for such an analysis do not yet exist for many of the issues that need to be weighed. Therefore, we have adopted a qualitative form of analysis that identifies issues in the

critical path of ATC evolution and specifies a framework for how quantitative data, when available, should be used to resolve those issues. Finally, we outline a program of empirical laboratory research which could provide those critical quantitative data.

Three alternative scenarios are presented in Section II. The first is a *Baseline* case which encompasses most of the current aviation-related research projects that are developing conservative technologies for ATC. The second scenario describes *AERA* (Automated En Route ATC), the FAA-sponsored R&D program to fully automate ATC functions.

The third scenario, which we term *Shared Control*, posits a number of automated aids which enable the human controller to retain ultimate control and still safely handle more aircraft. It is a technically more conservative scenario than AERA, with less lofty goals but more certain outcomes. (Two other systems that we do not consider to be viable options at this time are described in Appendixes A and B: a *Satellite-Based* ATC system proposed by various aerospace firms and a novel *Electronic Flight Rules* system in which sophisticated black boxes onboard individual aircraft would perform most ATC functions in a truly decentralized way.)

Since the Baseline scenario simply continues "business as usual," the analysis in Section III concentrates primarily on the AERA and Shared Control scenarios. Our conceptual framework for this analysis is presented in the form of four key principles:

- Cost effectiveness.
- Technical conservatism.
- Evolutionary progress.
- Human involvement.

Because these principles provide the foundation for the evaluation that follows, they should be weighed carefully against the reader's own axiomatic criteria for evaluation.

Sections IV through VI describe and contrast different aspects of the three alternative scenarios on the basis of our four key principles. Section IV compares the roles of the controller in each scenario; Section V compares the three concepts technically and economically; Section VI reviews the important differences among the scenarios and summarizes our recommendations to the FAA's research program. Specifically, we suggest that the AERA design be more liberally interpreted from a human-factors point of view, that the planned automation capabilities be scaled back to recognize the complexities inherent in this domain, and that the planned future role of the human ATC specialist be expanded rather than diminished.

II. ALTERNATIVE ATC SCENARIOS

A wide range of technological options exists for meeting future needs of the ATC system. Some of this technology has almost reached the deployment stage; some is outside the laboratory but still needs considerable engineering before deployment; and some is only conjecture from state-of-the-art research.

Our investigation will employ the concepts of *systems* and *scenarios*. The alternative ATC systems we discuss consist of numerous components (e.g., communication, surveillance, problem-solving, and management subsystems). The conjunction of these components forms a snapshot of a full ATC system, and the linking of the developmental phases of these systems over time comprises a scenario description. In other words, an ATC scenario emerges when we "string together" those interim ATC systems that might realistically form a progression from now to the turn of the century. Performing this synthesis repeatedly to accommodate many such coherent pathways produces alternative scenarios which can subsequently be evaluated.

The scenarios described below illustrate the wide disagreement in the ATC community over the best means of achieving the goals of increasing safety, making fuel-efficient routings available, and increasing controller productivity. The disagreements stem from widely different perceptions about what can be done and how to do it. Some observers, for example, are extremely optimistic about technological solutions, while others see policy-based solutions as more consistent with the nation's economic priorities. Within the four categories of scenarios to follow, we have tried to capture these varying perspectives about the future of ATC.

BASELINE

The Baseline case is a "default" scenario, in which the FAA simply continues to develop and deploy promising system components already under investigation. These include on/near-airport systems, surveillance-system improvements, ATC-facility improvements, and cockpit-based improvements.

On/Near-Airport Systems

Microwave Landing System (MLS). Allowing more numerous and more reliable approach paths to existing airports, MLS will presumably increase capacities somewhat. However, such increases are ultimately limited by minimum inter-arrival times over the runway threshold(s). MLSs may begin operating as early as 1985, but 1990 seems to be a more realistic time frame.

Wake Vortex Avoidance System (WVAS). Significant reductions in separation minima might be achieved with a system that can reliably report on vortex activity behind approaching aircraft. A successful WVAS could thereby significantly increase airport capacities. The time frame for WVAS is also the late 1980s, although technical problems make the implementation date uncertain.

Wind Shear Advisory System (WSA). Another limiting factor for aircraft approaching an airport is the presence of major wind changes close to the earth's surface. WSA will make a major contribution to safety rather than to increased capacity. It is also slated for installation during the mid-to-late 1980s.

Surveillance-System Improvements

Discrete Address Beacon System (DABS). Replacement of the current surveillance system with DABS has been studied for many years, and work on the system has advanced to the field-testing stage. DABS will improve radar coverage and reliability and will provide an air/ground/air datalink capability.

Collision-Avoidance System (CAS). Many collision-avoidance systems are under development. Some rely on ground-derived surveillance information; others are strictly cockpit-based and operate independently of ground radars. Some would automatically warn pilots of impending collisions in two or more stages (e.g., "proximity warning" followed by "alert"), and most would compute and recommend avoidance maneuvers for the aircraft involved. The first such system, T-CAS (Threat Alert and Collision-Avoidance System), is scheduled for installation by the end of 1984.

ATC-Facility Improvements

Replacements to the 9020 Computers (9020R). Current ATC computers were designed and built during the 1960s and early 1970s. They remain reasonably reliable and capable, but present load factors imply a need for them to be replaced no later than the late 1980s. New

architectural designs incorporating functional distribution should foster even higher reliability, as well as enabling much greater expandability and flexibility. In the Baseline scenario, 9020R software would be functionally equivalent to that in today's system.

Electronic Flight-Strip Displays (ETABS/TIDS). Flight-plan and other relevant flight-data information are given to controllers via flight strips, currently printed on strips of paper. Electronic displays of this information would increase productivity and enable more timely dissemination of flight-data information; demonstration programs of such displays already exist [1]. Electronic flight-strip displays could probably be fielded in conjunction with the 9020R system.

Flow-Control Automation. Current techniques of monitoring and controlling for delays at saturable airports could presumably benefit from increased levels of automation. Current FAA R&D plans include flow-control automation, but uncertainty about what will be done, and when, is still relatively high.

Cockpit-Based Improvements

Advanced Flight Management Systems (FMS). New Boeing 757/767 aircraft will contain over 100 microprocessors in cockpit automation, controlling almost every onboard system.¹ These new flight management systems will have precise four-dimensional navigation capabilities, enabling aircraft to be delivered over exact points in space at exact times.

Advanced Navigation Systems. Improved Loran and satellite navigation systems may begin displacing the nation's VORTAC system, although general aviation will undoubtedly rely heavily upon it for the foreseeable future. This change will affect the ATC system little, however, since microprocessor-based navigation systems already permit point-to-point routings in many aircraft.

Taken together, the above components constitute the least complex, least uncertain scenario we consider. Of course, if some or all of these developing systems do not emerge, even less capable ATC environments may result. Given historical and (especially) current sociopolitical trends, we do not consider this Baseline scenario to be especially optimistic or pessimistic; it is simply a harvesting of seeds already sown.

However, these technologies alone may not be able to meet the projected demand for ATC services, and demand-management tech-

¹Personal communication with Robert W. Sutton, Boeing Commercial Airplane Company, Seattle, Washington, 1980.

niques may be required. Whereas ATC has historically expanded as necessary to meet the unconstrained needs of airspace users, officials are now publicly suggesting that demand-controlled expansion of ATC may have to be halted—that allocation of services, rather than expansion of services, may be the watchword in a future “era of limits” [2]. According to H. Safeer, the FAA is actively considering such an alternative [3]:

... as the costs of expanding existing facilities and constructing new ones become increasingly prohibitive, more attention has been paid to alternate, low investment cost or noncapital-intensive techniques for accommodating increased demand.

These alternatives are generally of three types:

1. Alternative facilities to off-load congested airports (satellite, reliever airports);
2. Administrative (imposing maximum limits—quotas—on the number and type of operations which may use a specific airport or runway during a given time interval); and
3. Economic (charging variable landing fees, differentiated by time of day and by location; auctioning available landing and takeoff slots).

These last two measures do not physically expand capacity, but they can postpone the need for physical expansion by promoting more intensive and more economically efficient use of existing capacity.

Severe service shortfalls might bring even more restrictions, like those formulated in the FAA contingency plan for controller strikes [4]:

- Certain classes of flights, such as long-haul air carrier service, might be given precedence over others, such as general aviation.
- ATC might extend its reach even further into the pre-takeoff stages of flight, perhaps even to determining which aircraft are allowed into the system at all.
- Questionable services, such as flight following or radar-assisted sequencing of VFR flights, would be eliminated, redefining the advisory nature of current-day ATC.

Such radical changes may come to pass if the demand for air traffic services cannot be met in the future. And although this scenario postu-

lates increasing the size of the ATC service as demand grows, the Baseline systems may be self-limiting on a purely technical basis.

AERA

Technological advances in automated control permit consideration of quite a different scenario for the future. Each of the devices or systems mentioned above improves one aspect of the nation's ATC system, but ATC authority remains firmly in the eyes, ears, and minds of human beings poised over radar scopes. Suppose we could virtually replace these fallible human beings with a set of computer modules which could manipulate aircraft tracks so well that human intervention with individual aircraft would be necessary only in response to a major perturbation (e.g., a massive computer failure, or extensive storm-front passage). Suppose this computer system were able to automatically compute conflict-free clearances for aircraft under surveillance, to automatically transmit these clearances in a timely fashion, and to automatically monitor for compliance, taking corrective action as required.

The FAA is making exactly these suppositions in its AERA R&D project. The projected AERA *system* has been described in detail in a number of documents over the last few years [5,6], as well as in a recent position paper by a specially appointed panel of experts [7]. But no AERA *scenario* has yet emerged, so we have created one which faithfully represents the intentions of the research program and the systems that are to emerge from it. Our scenario is based on statements of the AERA designers and their published plans [6,7,8].

In the AERA scenario, computers would make *all* time-critical ATC decisions, at least for en route high and transition sectors. Responsibility for conflict recognition and resolution, as well as for flow control, would be officially transferred from the human controller to the machine. The human controller's role would be that of a "system manager" who ensures that the automation is performing its assigned functions properly and intervenes as required to handle exceptions.²

The technological goals of AERA are relatively straightforward. Figure 2.1 shows the major automated functions of AERA. The modules that perform these functions can be informally described as follows:

- *Surveillance/Flight-Plan Datalink.* Inputs and translates 9020 or 9020R information into a form usable by the other AERA modules.

²This concept of "human as manager" has caused much consternation within the ATC community, since everyone seems to have his own interpretation about just what such an AERA *system manager* should be doing. We will consider this issue at length later in this report.

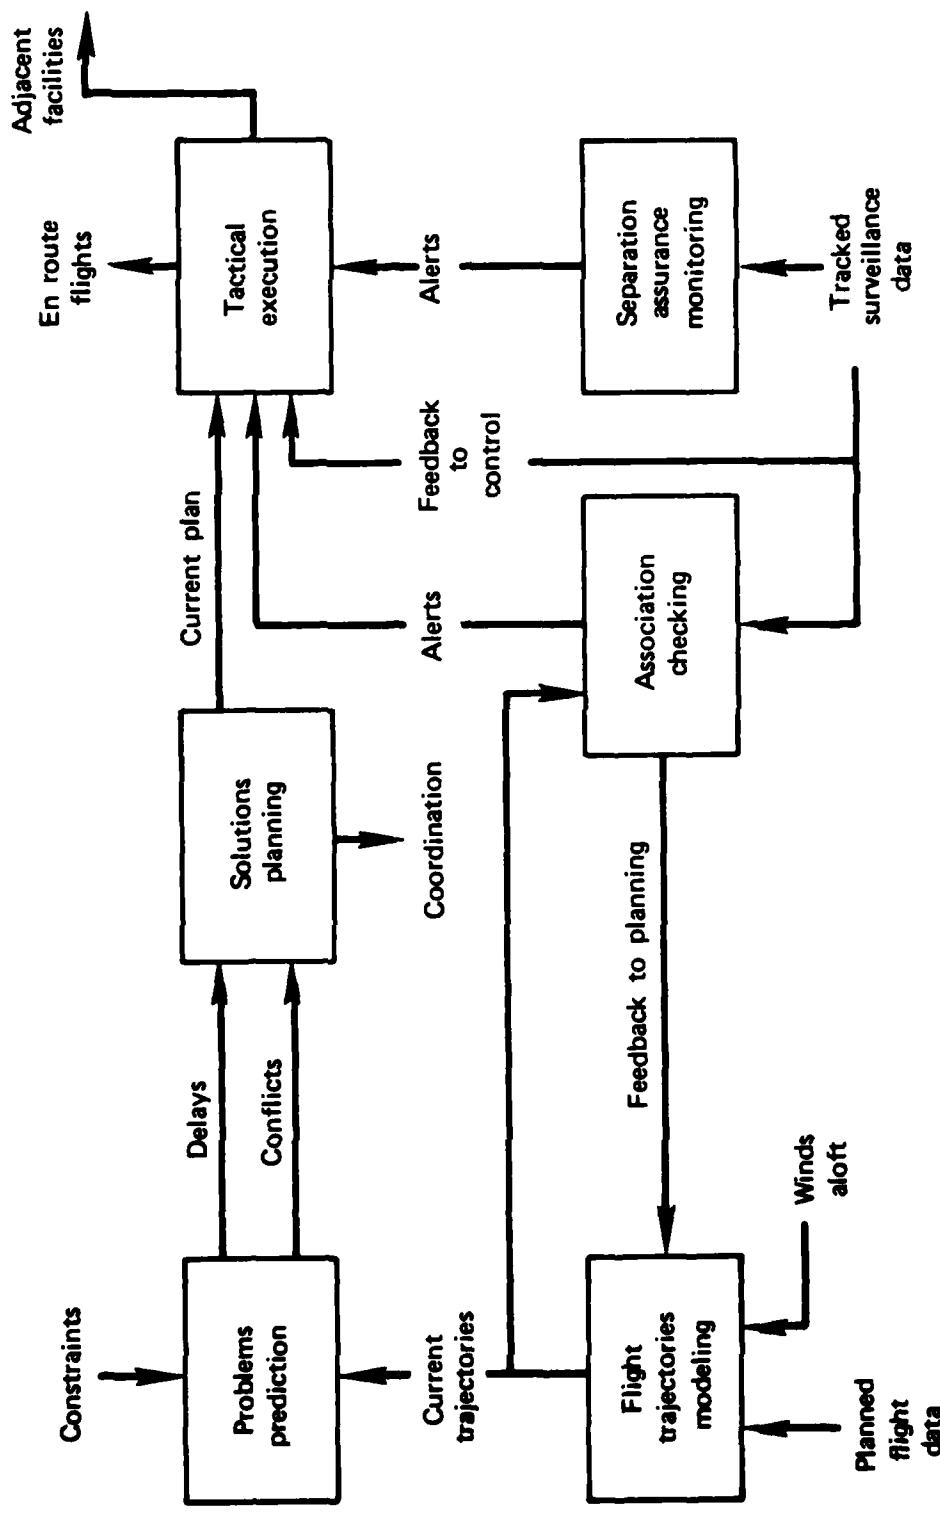


Fig. 2.1—Major AERA functions

- *Separation Assurance Monitor.* Performs tactical conflict monitoring and off-track monitoring. This module notifies the human controller whenever a potential problem is encountered and issues last-ditch resolution commands directly to the aircraft involved.
- *Strategic Planner.* Performs profile generation and strategic planning. The strategic planner would perform the longer-term decisionmaking task of deciding which aircraft are candidates for course revision, based on estimates of speed, heading, altitude change, converging courses, delay requirements, and the like. Its output would be the high-level instructions used by the tactical executor.
- *Tactical Executor.* Performs tactical command generation. This module set would translate high-level instructions, such as "Pass aircraft X behind aircraft Y," into the specific commands required for satisfaction of the implied goal.
- *Man/Machine Interface.* Provides an interface between controllers and the AERA problem-solving modules. Displays may roughly follow the design set forth in Ref. 8. Several options exist for partitioning the overall ATC task between controller and machine: The controller may be required to deliver by voice radio the clearances "suggested by" the automation; he may be required to approve such clearances prior to delivery by simulated datalink; or he may have only veto power over such clearances.
- *Failure Recognizers and Reconfigurers.* Several schemes for recovery have been postulated, all based on the premise that if the human manager is not routinely in the control loop, he cannot react to system failures quickly enough to be effective. AERA is designed with redundant, fail-safe processors to guard against complete hardware failure. It uses multiple layers of separation-assurance software, so that subsystems are continuously checking each other for potential conflict situations. If a catastrophic centerwide failure should occur, the center's AERA will activate backup clearances and initiate a stabilizing process which, depending on the specifics of the failure, will divert its traffic to adjacent centers or initiate manual control procedures locally.

The FAA plans extensive testing of AERA before deploying it. After a laboratory development phase during the early 1980s, the system will be tested at the FAA Technical Center in Atlantic City and then in a real ATC center. Interfaces with either the 9020s or their replacements

(if available by then) will be constructed so that an AERA prototype can be presented with live data and real situations. Only after the system has proved itself repeatedly in such "shadow mode" operations will it be deployed in centers as the primary controlling entity.

However, some components of the AERA concept, notably those involving automatic planning of navigation-direct and fuel-efficient profiles, may be fielded earlier as controller aids. These aids, which will probably be installed during the late 1980s, should allow controllers to authorize more direct routings by using simplified AERA algorithms to compute conflicting situations. But AERA advocates repeatedly point out that these components will achieve the hoped-for gains in productivity only when a complete system is available. In keeping with this philosophy, our scenario does not envision a significant AERA impact before the end of the 1980s.

Optimists believe an early AERA could come on-line at a real ATC center around 1990. Pessimists suggest 2000 as the earliest possible date. Our scenario takes a moderate position on the timing of its implementation: During the early 1990s, full-scale testing of AERA I is completed and a contract for construction is awarded; by the mid-1990s, some version of AERA will be on-line at some centers; and by the late 1990s, AERA should be on-line and "in control" at all centers.

According to current plans, its contributions at that time will be confined to high-altitude and transition airspace sectors. Although controller staffing levels in terminal areas will continue to climb, en route centers will experience first a leveling and then a decline in manning levels as AERA takes over the routine en route ATC functions. Table 2.1 summarizes the major events in this scenario.

SHARED CONTROL

The AERA scenario raises uncertainties that make it a very high-risk proposition: Is the role of "system manager" viable? Can automation indeed handle almost all traffic situations with no human intervention? Will the automatic error-detection and reconfiguration procedures work? Can the touted gains really be achieved? Considering these uncertainties, is there an acceptable alternative—that is, one that meets the projected demand for ATC services but relies less heavily on untested automation?

Our answer is, Maybe. We have postulated a scenario based on such an alternative. This alternative, Shared Control, parallels AERA in many respects but focuses on keeping the human in the control loop at

Table 2.1
SYNOPSIS OF AERA SCENARIO MAIN EVENTS

1981:

- Development continues in laboratory on highly automated control algorithms and human interfaces.
- Evolutionary deployment plan is developed.

1985:

- AERA testbed is complete; laboratory experimentation demonstrates feasibility of giving machine primary separation assurance responsibility.
- AERA failure modes are defined and design work completed.
- Contracts are awarded for initial AERA modules which will provide controllers with automatic profile-generation and conflict-checking aids.

1990:

- Replacements to 9020 computers come on-line at all centers.
- AERA testbed field tests are complete; contracts are awarded for construction of first fieldable AERA system.
- First AERA-derived automated aids are fielded and used in en route centers.

1995:

- AERA is on-line at one center for extensive testing.
- All centers are using some set of AERA-derived automated aids.

2000:

- All centers have AERA on-line.

all times. This concept arises from analyses which suggest that man is likely to be a poor system monitor unless he is actively involved in the control process [9]. It continues the evolutionary development and deployment of automated aids—not replacements—for air traffic controllers through the next two decades. The Shared Control scenario reaches much the same level of automation as AERA by the turn of the century, but the pathway there is markedly different.

In this scenario, the controller's verbal workload will initially be reduced. By about the mid-1980s, DABS, ETABS, and other digital communication support devices will be integrated to enable a significant portion of air/ground communication to be made digitally. Special cockpit "digicoms" will enable pilots to send and receive encoded messages. For flights that have no digicom capability, a voice generator might transform the controller's digital commands into "spoken" clearances transmitted over the usual VHF communication channels.

The controller's mental workload will be reduced by a special digi-com interface called a *Tactical Communications Manager* (TCM). The TCM provides an electronic blackboard upon which controllers can post clearances for later delivery according to anticipated temporal or spatial conditions. For example, instead of having to recall a planned off-airways vector or anticipated altitude change, a controller will enter a planned command into the TCM for issuance when the appropriate airspace pattern develops. Libraries of standard procedures will facilitate the entry of complex but frequently used plans. The TCM should function as a notepad, assisting the controller in memory functions and freeing him to concentrate on planning for the future.

Since clearances will be routinely stored in the computer if the TCM is being used properly, a monitoring and planning aid can be added which utilizes these clearances to predict the future. The first of these aids, a *Plan-Ahead Monitor* (PAM), is similar to but simpler than AERA's Strategic Planner. It should aid the controller's visualization process, back up his separation-assurance control function, and free him of the need to perform track, conflict, and flow prediction mentally. (Some rudimentary aids exist for the latter function even today.) PAM is designed to use stored aircraft performance parameters, airspace knowledge, and planned clearances to dynamically display potential "futures" on controller command. It will have numerous modes of operation:

- A *background mode*, which performs global conflict monitoring and alerting continuously. This mode can be thought of as an "intelligent" version of the current-day conflict alerter, in that intended flight-path alterations will be known to PAM through the stored digital clearances. It implements the functions planned for AERA's Separation Assurance Monitor.
- A *time-based look-ahead mode*, in which time can be manipulated according to controller directives input via an appropriate analogue device such as the current trackball. In one such mode, spinning the trackball quickly to the right would advance time quickly forward, and an auxiliary planning display would show aircraft moving "supersonically" across the screen. Spinning more slowly might cause time to move more slowly—in the vicinity of some future interesting event, for instance. An aircraft-specific mode might also be available in which a flight path could be artificially cursor-controlled and a clearance plan automatically generated in response to that motion. In this mode, the controller can place the planning screen's cursor over the subject aircraft and "maneuver" it in fast time to achieve

some desired profile. Other aircraft on the screen will be updated according to their velocities relative to the subject aircraft, so that the screen always shows a consistent picture of projected futures. When the profile is completed, PAM will automatically construct and store as a clearance plan the required vector/altitude/speed commands to effect that profile.

- A *spatial look-ahead mode*. The future need not be presented in a time-varying fashion; instead, aircraft profiles might be drawn and conflicts shown directly over the space in which they might occur. Either horizontal or vertical profiles might be used in this presentational concept.

A rudimentary simulation of the space-based look-ahead mode is already being demonstrated in the laboratory, and present-day simulation techniques would be suitable for PAM's software. PAM will enable the controller to take advantage of the reduction in monitoring workload achieved by the TCM by providing a more precise picture of the future than he can project mentally.

Expansion of PAM beyond simple look-ahead to include a modest planning function characterizes the mid-1990s stage of this scenario. A set of planning aids, which we call *Autoclear*, will be deployed which roughly parallel AERA's Strategic Planner and Tactical Executor. These aids are a straightforward extension of PAM's aircraft-based look-ahead mode described above. Individually invokable modules for most planning functions will be available to the controller at the push of a button. He might request advice on strategic options for a particular aircraft or group of aircraft. He could send this strategic plan, a modification of it, or a new one of his own design to a tactical executor which will then create a specific sequence of commands for issuance by the TCM. As an integral part of a man/machine system for generating clearances, the human controller will generally reserve the higher-level decisions for himself and use Autoclear to perform the low-level details. He should rarely be involved in the minute-to-minute operations of conflict monitoring and clearance issuance, but should be able to spend even more time than before designing efficient yet safe routes and flow patterns.

To confront problems in the increasingly congested terminal areas expected during the mid-1990s, this scenario emphasizes the development of "intelligent" cockpit displays of traffic information (CDTI). Onboard processors will use DABS-transmitted surveillance data to electronically inform pilots of the local traffic [10,11], and the "intelligent" CDTI of 1995 will efficiently filter out irrelevant traffic and interact with various collision-avoidance systems when a conflict does occur.

Finally, by 2000, the Shared Control scenario posits two variations of Autoclear: a relatively simple version adapted for use at the terminals and a more complex one, called *Autoclear II*, for the en route centers. The "old" Autoclear will consist of numerous distinct modules which, by the late 1990s, will have undergone many revisions, updatings, and enhancements based on feedback from its users. Autoclear II will integrate all of these modules under the control of an *Executive* problem-solving system. The Executive will then monitor the state of each sector environment, automatically activating appropriate individual functions. A complete flight profile and its attendant clearances will be constructed, transmitted, and verified by the Executive as required.

Although Autoclear II will have roughly the same capabilities as are planned for AERA, it will typically not be allowed to manage a sector alone. Autoclear II will perform the lowest-level separation management functions and will be used to construct and issue "reasonable" clearances at high traffic densities. But the human controller, with his superior global perspectives and situation-specific knowledge, will frequently override the Executive and manipulate the Autoclear subfunctions directly to produce customized—and better—clearances than the machine would. In so doing, he will continue to perform many of the ATC tasks he does today. The difference is that in 2000, he will rely heavily upon a vast array of automated aids.

Figure 2.2 illustrates the automated functions we envision for this scenario, and Table 2.2 details its developmental time schedule.

OTHER SCENARIOS: HIGH-TECHNOLOGY ATC

Solutions that rely on even more advanced technologies are also being discussed within the ATC community. We describe two such scenarios in Appendixes A and B. The first uses satellite-based communications to coordinate a nationally centralized ATC system. This extraordinarily complex system demands extensive development of new technologies and replacement of all existing facilities. We conclude that it is too revolutionary and too costly for its uncertain benefits. The second high-technology scenario, termed Electronic Flight Rules (EFR), uses onboard processing to shift all separation responsibility to the cockpit. This scenario also requires revolutionary technological development and is, we feel, likely to be significantly less safe than either the Baseline, AERA, or Shared Control scenarios. We shall not discuss these high-technology scenarios further in this analysis.

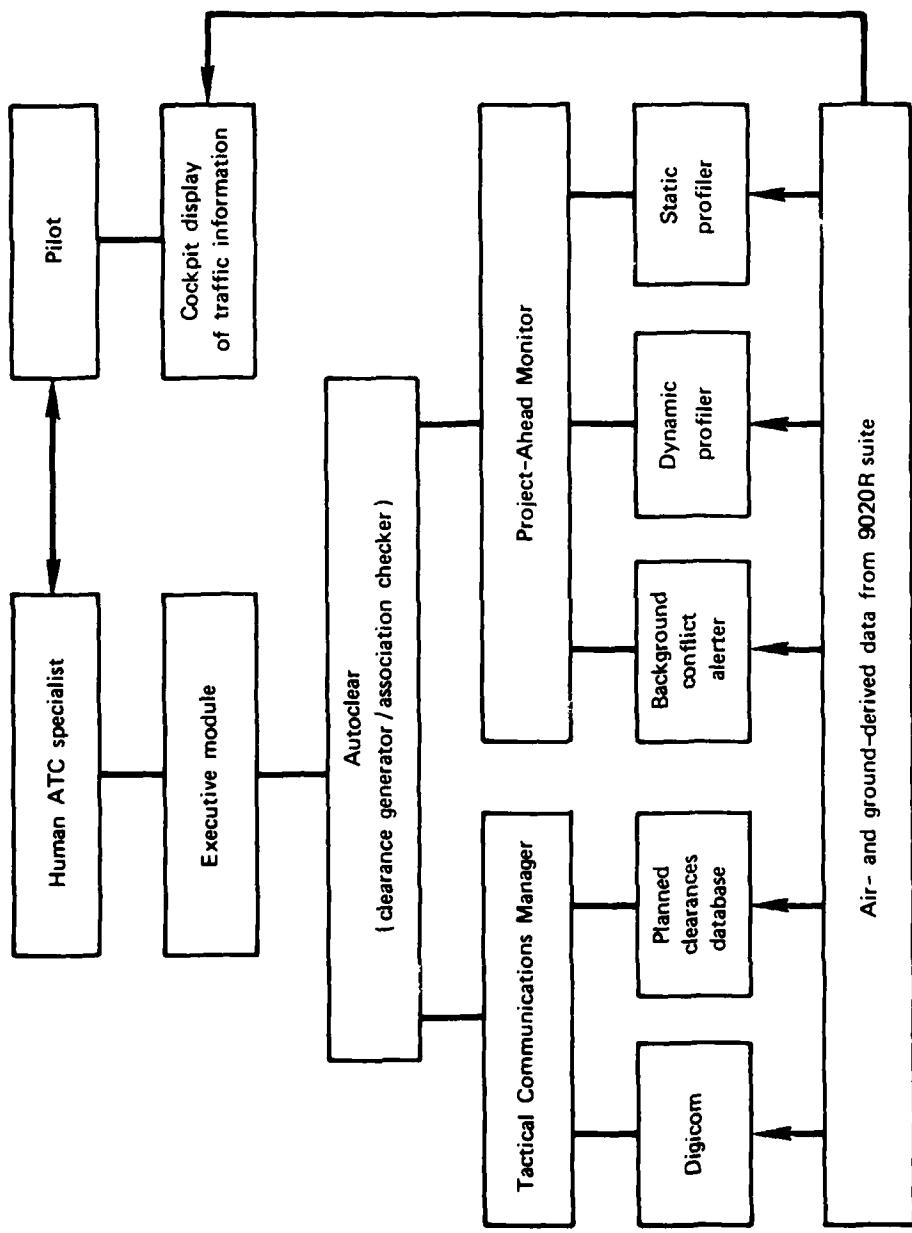


Fig. 2.2—Shared Control final system module configuration

Table 2.2**SYNOPSIS OF SHARED CONTROL SCENARIO MAIN EVENTS**

1981:

- Policy decision is made to emphasize AERA functional modules as individually invokable controller aids.
- Digital communication requirements are defined; air ground protocols are designed to support these requirements.
- TCM testbed is under construction.
- PAM designs are under development.

1985:

- DABS is operational in selected terminal areas; digicom is being field-tested; fleet equipage is still low but increasing rapidly.
- TCM and ETABS are implemented in all centers.
- PAM is undergoing field tests at selected centers.
- Initial Autoclear functions are defined; initial designs are completed.

1990:

- PAM is on-line in all centers.
- Some Autoclear functions are in field tests; some are still in laboratory.

1995:

- Autoclear is on-line in all centers.
- Refinements suggested by controllers result in new releases of various functions from time to time.
- Executive to unify Autoclear functions is defined and undergoing laboratory testing.

2000:

- Autoclear II, including Executive, is on-line in all centers.
- Evolution continues as refinements increase Autoclear II performance.

III. GUIDING PRINCIPLES

The scenarios described above are designed to achieve similar goals by extremely diverse means. They present many options that must be carefully weighed and analyzed before choices are made that will affect R&D programs costing millions of dollars.

The FAA has already asked a number of organizations to address the question, What should the future of the ATC system be? These organizations have used a variety of methodologies—eliciting and statistically analyzing the opinions of panels of experts [9,12], running demonstration projects using highly simplified domains [5], creating designs from "scratch" [13], and getting the ATC controllers themselves to assess their future needs [11].

This problem encompasses many issues common to the definition and design of human/machine decisionmaking systems. We have thus approached it from the perspectives of computer science, engineering, human-factors psychology, and the emerging field of cognitive science. We have attempted to use existing quantitative data—projected growth rates of the controller force under various conditions of automation development, projected demand for ATC services, and previously successful applications of related technologies—but most of these have proven to be of questionable relevance to our analysis.

Consequently, we have relied heavily on our own observations of current ATC operations and laboratory simulations of ATC tasks. These observations led us to adopt four qualitative principles for evaluating potential systems and scenarios: cost effectiveness, technical conservatism, evolutionary progress, and human involvement. Each principle embodies our value judgments and is viewed by us as axiomatic; by rejecting one or more of these principles, one can logically derive conclusions about the possible systems and scenarios that are different from those presented here. However, these are the principles that have emerged consistently both in our research and in the research of others.

COST EFFECTIVENESS

Each scenario provides some improvement in ATC system safety, aircraft fuel-efficiency, and controller productivity for a given expenditure of money. Some scenarios may net more improvements than others for the same expense, or they may net the same level of improvement earlier for the same cost. Some scenarios may require a certain mini-

mum investment before paying off at all. Some may require massive expenditures of R&D dollars before we can even know whether or not they might pay off.

Our principle of cost effectiveness gives the highest ratings to those scenarios that have the earliest, biggest, or most certain payoffs and the widest application. A scenario that achieves modest gains quickly may be more cost effective than one with a greater but delayed payoff. Similarly, a scenario that is almost certain to succeed is rated above one that costs about the same but whose outcome is less certain. Since no reliable cost and performance data are currently available, we can offer only qualitative estimates of the costs and risks of pursuing any of the scenarios.

TECHNICAL CONSERVATISM

The future ATC system must be built upon a foundation of reliable, expectable, conservative technology. Unfortunately, there is no universally agreed-upon definition of what constitutes a "conservative" technology. Some will argue that only today's proven technology is conservative, even though our horizon extends 20 years beyond these concepts. Others will argue that 20 years is a long time, that today's emerging technology will be well-established by then and thus can be regarded as conservative for planning purposes. Still others will suggest that anything that can be imagined within this time period should be included.

We favor the moderate position. Although hardware capable of supporting a highly automated, high-performance ATC system is likely to become quite reliable during the next 20 years, and software to perform most routine controller tasks will also advance significantly, the ATC system cannot gamble on these expectations. It must design and develop scenarios within the context of capabilities that can be convincingly demonstrated in today's laboratories. We must remember that the potential costs of R&D failure or delay in a highly interdependent ATC system are very high. Projected performance gains must be weighed very carefully against those costs for worst-case outcomes.

EVOLUTIONARY PROGRESS

The principle of evolutionary progress covers two important phases of a scenario: its deployment and its development. During deployment, each succeeding system that is introduced as a scenario progresses

must be smoothly incorporated into the existing ATC environment, perturbing ongoing operations as little as possible. Clearly, human users must become familiar with each new system before abandoning the procedures or facilities it replaces.

Similarly, graceful evolution is important during system development. No advanced system will spring from its designers' minds fully matured; time and opportunity must allow the system's users to communicate their changing needs to the designers. As users adapt to the new system, a feedback pathway must exist for them to suggest changes they could not have anticipated before using initial versions of that system, to say to the designers, "Now that I've had some experience with what you've given me, I know what I *really* needed in the first place."

HUMAN INVOLVEMENT

The principle of human involvement is surely the most controversial and, for us, the most crucial. It asserts that the human role should not be determined solely by what the machine can do best, but also by what the human must do at all times in order to support or maintain his performance for those tasks the machine cannot or will not be allowed to do. This principle is usually taken to mean that the human must be continuously and intimately in the control loop. It does not become controversial until we try to get consensus on just what that means. Some agree with the following assertions from the FAA's Tiger Team on AERA [7]:

The controller is in the loop in today's system. The controller is not in the loop in AERA with respect to aircraft control, neither is he required to monitor clearances. AERA or a pilot might ask the controller to monitor or handle certain situations, but this is control by exception. The controller is the manager of AERA and traffic flow, but does not control individual aircraft. He is provided with system status, weather, traffic demand and capacity displays to perform his managerial responsibilities, as described below.

In this manner, the controller is relieved of routine, which should minimize errors, and has the more rewarding responsibility of creatively using ATC and AERA assets to satisfy traffic demand.

Others will agree with S. Poritzky, director of the FAA's Office of Systems Engineering Management [14] :

We have talked about the controller as "system manager," but other than that it sounds nice, we don't know what a system manager is in this context. Does he simply look at a panel of red and green lights, and start to worry when the light turns from green to red? Does he actually perform the same process that the machine is performing so that he can take over in case the machine fails? Just what does a system manager do? In an automated system, who has the final responsibility for separation? If a catastrophe should occur, who is responsible—the controller, the machine, the computer programmer, who? If you let the automatic process operate in such a way that in the event of a failure, the human controller can take over the whole show instantaneously, then why bother to do it at all? If the process normally runs automatically, but fails occasionally, leaving the job to the controller, how does he maintain proficiency in what is—by common consent—a tough control problem?

Researchers in Great Britain take this sentiment even further [15]:

It is argued here that this primary involvement of the controller is a sine qua non for computer-based ATC systems. Without it, there is a danger of the controller's becoming remote from the practical situation . . . and of being less than efficient in intervening when the need arises. This principle is at variance with some current U.S. systems research in ATC [here some AERA work is cited], which proposes that new conflict-free clearances be generated automatically by program and presented to the controller for a check before being delivered automatically to the aircraft. Hard evidence is notoriously difficult to obtain on such issues, but there would seem to be a risk of the controller's losing contact with the traffic situation as a result of such a passive role. . . .

Even if we keep the human "in the loop," however, we must also agree with Poritzky's statement at a recent Office of Technology Assessment seminar [16] :

If we are to provide a high level of flexibility in aircraft operations and permit conservative fuel use, we believe the task of air traffic control—that of traffic separation and efficient flow management—will require juggling more variables than can be done successfully by human controller teams alone.

So the question of how much and what kind of human involvement is necessary to satisfy this principle stands at the very heart of our analysis. More than any other principle or metric, human involvement

may be decisive in choosing a "best" scenario. Each scenario presents a markedly different option: The Baseline case keeps the human firmly and completely in the loop; the AERA case takes him out of it; the Shared Control case keeps him somewhere in the middle during most of its span, although it begins to converge with AERA in its latter stages.

The answer to this question must be based on an intimate understanding of human and machine capabilities and limitations. If one can convincingly argue that the whole of ATC can be automated, then how or why the human being would fit into the system becomes merely a political issue. If ATC cannot be completely automated, then allocating tasks between man and machine requires balancing human and machine skills optimally. This problem is not unique to ATC, of course, and there is a wealth of human-factors and system-design studies on the subject. Mertes and Jenney [9] have compiled and summarized the following important conclusions from these past studies of human and machine performance characteristics:

- Man is an unreliable monitor. The more passive his role in a system the more he tends to withdraw from the system by letting his attention wander or even by going to sleep. If it is desirable that man serve as an emergency backup, then he should be given tasks to keep him aware of what is happening in the system so that he can take over when needed. It may be necessary to give him these tasks even though they could better be done by a machine. (p. A-3)
- The human operator should not be assigned monitoring tasks that require continuous attention to a display unless absolutely necessary. (p. A-3)
- Humans are relatively poor, with respect to machines, for performing routine, repetitive tasks. (p. A-3)
- In perception the human has distinct advantages over machines. Humans perceive patterns, not isolated bits. These patterns are not restricted to one sensory modality, but may include some or all of them. . . . Man can also perceive patterns of events occurring over time and thereby anticipate events; this is behind much of his ability to learn. (p. A-18)
- The ability to reason inductively, that is, to make generalizations from specific observations is perhaps man's greatest claim to fame. . . . He is the only available computer able to solve problems by logical induction. (p. A-24)

- Generalized information processing and decisionmaking should be performed by personnel where:
 - a. Pattern perception is important (especially where patterns may change in size, position, or energy configuration (types and strength levels) under different conditions).
 - b. Long-term storage of information is required.
 - c. Insight, discovery, or heuristic problem-solving is required.
 - d. Decisionmaking and learning in a complex changing situation are required.
 - e. Ability to improvise and adopt flexible procedures is important and, within the state of the art, cannot be built into a machine program.
 - f. Number of low-probability events which might occur is high and the cost or capacity of machine programming is exceeded by the requirement.
 - g. Inductive reasoning is required, i.e., a requirement exists for generalizations to be made from the specific events. (p. A-11)

The implications of this synopsis of human-factors literature stand out clearly: Routine, repetitive operations should be automated if possible, but the handling of exceptional and "fuzzy" information requires man's intellect. If his intervention is to be required, then he must be assigned a suitable level of task involvement to keep him attentive and ready to perform his duties.

Two controversial issues prevent these conclusions from generating a consensus about human involvement in a highly automated ATC system: The first is the issue of exactly how complex and "fuzzy" ATC problem-solving is—that is, how much of it really requires human capabilities. The second is a disagreement about what "suitable level of task involvement" means.

Most observers of ATC operations concede that much of what controllers do is routine, repetitive, and automatable. Some observers go beyond that by asserting that almost all of the task is automatable, that the complex pattern-matching and decisionmaking behaviors which characterize human performance in ATC are really just poor approximations of mathematical projections which a computer can perform much better. However, our observations at local centers and TRACONs have convinced us that much, if not most, of a controller's time is spent on tasks that require distinctly human skills: negotiating flight-plan

changes with pilots, vectoring aircraft around rapidly changing severe weather, deciding upon general operational configurations with other controllers, and the like. These tasks also require experience, maturity, and flexibility—the blips on those screens are, after all, real people who change their minds and make mistakes. Clearly, human involvement in ATC means comprehending and responding to situations whose complexity mainly stems not from profile projections but from the breadth of man's shared experiences.

The "suitable level of task involvement" required to perform this complex role remains an open question. Some systems designers assert that a future human "system manager" does not need to continuously monitor and manipulate the individual aircraft on his screens, but that he should only have to handle more abstracted information, such as aggregates of planes or flow patterns. However, the "holistic knowing" which comes from an intimate involvement in every detail of the traffic control process may be necessary to sustain the complex controller behaviors mentioned above. In many instances—particularly in those that involve life-or-death situations—there simply may not be time to query the computer for an answer.

To the extent that humans simply back up the automated system—that is, exercise little or no control unless the machine functions unsatisfactorily—the resulting boredom of the task can present a safety hazard. Thackray, Bailey, and Touchstone [17] found that increasing the boredom of subjects monitoring radar displays increased fatigue, irritability, strain, and response times and decreased attentiveness and arousal. Thus, removing the responsibilities of controllers may lead to seriously deficient performance in situations where human intervention is required.

In summary, then, the principle of human involvement means that a future ATC specialist must be given enough automated assistance to enable him to manage increased traffic loads while still retaining enough control responsibility and information to manage the overall system operation. It means that the automation may assist, but not completely do away with, a controller task unless it can perform the task completely and reliably, as well as all the other tasks that "depend" on that task. The principle of technical conservatism further constrains automated performance of human tasks, leaving us with rather strict requirements for the human role in any automated ATC system.

IV. HUMAN ROLES IN EACH SCENARIO

Each of our three scenarios posits a different role for the future "ATC specialist." This section discusses those alternative roles.

We must initially distinguish the multiple human roles in any ATC system. Management personnel, data-systems specialists, radar-control teams, and trainees all contribute to the successful operation of an ATC facility. These role distinctions will endure. However, we shall limit our discussion to the sector control teams¹ per se, those individuals who man the radar sector positions, communicate with the aircraft, and control their passage through the sector.

To be sure, future control teams may differ in size and function from the present ones. An AERA control team, for example, may oversee much more airspace than is encompassed by a current-day sector. There may be a much larger team of information-processing specialists to tend the extensive automated systems. These specialists will be instrumental in recovering or reconfiguring during failures, but their routine functions will consist primarily of specializing and improving the facility's hardware and software.

Our analysis of human roles is strictly limited to the active control functions, which range from tight, open-loop manual control in the Baseline scenario to general, closed-loop managerial duties under AERA.

BASELINE

Except for an increase in coordination activities, the controller's role in the Baseline scenario will presumably be an extension of what he does today. By 2000, there will be twice as many controllers handling twice as many aircraft, and they will have a few additional automated aids such as ETABS and more reliable computers with the new 9020Rs. Average sector sizes will be smaller to keep individual control team loadings manageable, but basically, ATC in 2000 will parallel ATC in 1981.

We see a singularly difficult problem arising from this straight-line extrapolation process of adding more controllers who simply coordinate with each other more than at present: How small can sectors get before

¹A control team generally consists of one to three persons.

the problems of coordination overwhelm the control teams? That situation already exists at the Los Angeles TRACON "Downey" sector, where traffic loads regularly surpass a desirable maximum for a single sector. All attempts to split the sector have failed because of coordination problems between the newly created sectors.² We expect more such situations to arise as veritable atomic limits are reached on sector size. Coordination procedures already account for roughly half of a control team's workload, and as these procedures become more frequent, the time available for monitoring and planning the controlled airspace decreases, finally resulting in a sector which cannot be effectively controlled.

AERA

Installation of the first AERA system will herald a significant, radical change in the role of the ATC specialist. Instead of controlling individual aircraft as he does today, he will manage a massive automated system which will control the aircraft for him, under his direct supervision.

Human roles in the AERA scenario will be characterized by two distinct phases. The first phase will be an interim period while AERA is being introduced. A few partial-AERA control aids (like the planning aids identified earlier) will be provided, after which the full AERA system will be installed but will operate in a "background" mode while being configured to the particular center's airspace. In the second phase, AERA will assume full primary control and the specialist will cease to be in the control loop.

During the transition phase, a difficult transfer-of-control problem will face ATC decisionmakers. The specialist will still be responsible for controlling aircraft, yet the machine systems will have to be gradually given more and more of this responsibility. The machine will be making recommendations to the specialist which he cannot verify directly and which must therefore be completely correct and trustworthy. If the specialist chooses to accept inadequate recommendations, is he to blame for not overriding the machine? And if he rejects superior recommendations, thereby wasting an aircraft's time and fuel, has he also acted improperly?

One way to circumvent this dilemma is to build machine functions that are trustworthy and complete in their problem-solving skills before they are fielded. (We shall assume for the purpose of discussion that

²Personal communication with David Ross, Los Angeles TRACON Data Systems Services Officer.

this can be done.) The temptation is strong to then vest full and complete responsibility in the machine for those tasks that it has proven it can handle. An almost instantaneous transition to AERA or some significant subset of it may leave controllers unable to cope with their new managerial responsibilities, never having had a chance to get used to their much-altered new role.

This concept of "managerial responsibilities" also merits closer scrutiny. If the transition to AERA is in fact made successfully, the ATC specialist should become its system manager. But little information exists about precisely what an AERA system manager would routinely do. Is the human specialist to be left with obsolete skills and a few "fill-in" duties, such as voicing machine-generated clearances over the VHF radio and inputting pilot replies and requests? Such a role is outlined in Ref. 7 (Section VI), but more work needs to be done and the role remains poorly defined.

We have attempted to clarify this role by extracting the human functions specified or implied in the AERA design document and detailing them on the basis of our experience with other man/machine decisionmaking systems. A synopsis of these roles is given in Table 4.1.

In this listing of behavior patterns, the ATC specialist's routine role is that of a system monitor and special-case resolver. He will assign machine resources to ensure their efficient use and monitor the general system health. He will initiate failure-mode reconfiguration procedures if his monitoring turns up too many anomalies. He is expected to monitor aircraft tracks for suboptimal or erroneous machine handling and intervene appropriately to correct or improve the situation. (This function will be done by spot checks or in response to machine requests, since routine traffic loads will exceed human capacity.) He will revise machine-generated clearances to accommodate pilot requests, weather, and other special situations that the machine cannot handle, either because it is not programmed to perform that function at all or because its capabilities are inadequate for the situation. He will fine-tune machine problem-solving functions to meet special cases in his sector.

The prospects for this role definition becoming reality depend primarily on the capabilities of the automated control system. If the machine routinely handles virtually all of the traffic situations completely, leaving the human with little to do, this role is indeed manageable (although not particularly desirable or interesting). The main problems facing the human would be skill loss over time and lapses in attention caused by the low frequency of important events. These problems may be alleviated by frequent training and a requirement for regular reporting behavior when working a shift.

But suppose the machine cannot perform flawlessly and must ask

Table 4.1
HUMAN ROLES IN AERA

Control Function	Reason for Controller Attention	Controller Activities	Controller Processes and Tools
Routine Operations			
Monitoring AERA strategic planner	Maintenance of skill and vigilance in low traffic periods To note abnormally high aircraft densities Planning diagnostics are displayed by AERA	Examine AERA solution profiles Examine distribution of processing load across AERA modules; examine other AERA status data Request AERA to display system status, current plans, and planning constraints	Monitor plan view and textual displays Input requests for system status data View simulation output on planning display strategic planner
Monitoring airspace availability			
	Weather changes Changes in airspace restrictions Changes in operations at adjacent TRACON Changes in traffic load and mix	Receive or initiate communication with weather unit, pilots, or neighboring center Enter new airspace utilization options into AERA Direct AERA to simulate effects of parameter change View simulation output	Communicate via telephone Use command syntax and/or graphics interface to enter airspace modifications Input new airspace utilization policies

Table 4.1—continued

Control Function	Reason for Controller Attention	Controller Activities	Controller Processes and Tools
Handling pilot requests for flight-plan alterations	CDTI/DABS-equipped pilots request controller assistance	Receive communication from pilot	Communicate via radio
	DABS-equipped pilots request controller assistance	Use AERA to evaluate feasibility of desired change or to produce a new plan	Enter pilot request to AERA
	Non-equipped pilots request controller assistance	Communicate plan to pilot	Communicate plan and conditions to pilot
Monitoring flow control	Controller receives requests from adjacent non-AERA sectors for flow restrictions	Communicate airspace conditions and maneuver constraints to CDTI-equipped pilot	
		Receive communication from neighboring sector	Communicate via telephone
		Enter new constraints into AERA conditionally	Enter tentative revisions to parameters, using command syntax
		Negotiate restrictions with sector module on airspace availability	Direct AERA to simulate effects
		Enter constraints into AERA	Enter agreed-upon revisions
		Inform adjacent sectors of AERA-generated restrictions	

Table 4.1—continued

Control Function	Reason for Controller Attention	Controller Activities	Controller Processes and Tools
Monitoring sectorwide replanning	Severe weather Changes in terminal operation status	Examine AERA backup plans Examine distribution of processing load across AERA modules	Communicate via telephone Coordinate new routing requirements with other ATC facilities via telephone (e.g., receive ok to reroute aircraft; request flow restrictions) Direct AERA to simulate and display effects of plan changes
Regulating demand on and utilization of AERA capacity	To maintain efficient usage of AERA resources To prevent system saturation	Monitor system statistics (e.g., processor and communication utilization) If workload too high, adjust system parameters and/or negotiate restrictions with adjacent ATC and/or activate peak-load processors	Check status displays Hypothesize parameter changes to reduce workload Enter parameter changes into AERA Communicate via radio with adjacent ATC
Maintaining AERA data base	To receive data from non-AERA ATC (flight plans, hand-offs, flow restrictions) To control aircraft with no or failed transponder AERA alert due to absence of flight data for target aircraft	Receive communication from adjacent ATC Receive communication from aircraft Note AERA alert	Communicate via radio or telephone Enter new information into AERA using touch-panel display and/or keyboard Verify AERA acceptance of data and accuracy of data

Table 4.1—continued

Control Function	Reason for Controller Attention	Controller Activities	Controller Processes and Tools
Transmitting clearances and advisories to aircraft	Aircraft not DABS-equipped	Obtain information from AERA	Request tactical executor output for aircraft
	DABS fails for aircraft	Communicate information to aircraft	View output
	Non-DABS communication		Communicate via radio
Separating and metering individual aircraft	Aircraft require special services for which AERA is not designed	Override AERA conflict resolution and tactical execution modules	Receive flight plan Recognize special conditions
	Maintenance of skill under low traffic conditions	Control tactical execution of strategic plan	Negotiate a tactical plan with pilot via radio
Monitoring aircraft activity			Enter special constraints to AERA, using command syntax, planning display
			Monitor aircraft activity
			Communicate clearances via radio

Table 4.1—continued

Control Function	Reason for Controller Attention	Controller Activities	Controller Processes and Tools
Failure-Mode Operations			
Verifying voice radio frequency changes	AERA verification system fails	Receive communication from aircraft Acknowledge that aircraft is readable and controlled by him Acknowledge via radio	Receive radio call from aircraft Verify aircraft position on PVD Acknowledge via radio
Monitoring software recovery functions	AERA signals software error Controller notices incorrect AERA functioning	Continuously monitor system status displays (e.g., processor and communication utilization) Monitor switch of system into backup clearance mode Coordinate with adjacent sectors	Detect abnormal system status Use AERA MMI to coordinate with still-functioning AERA modules Communicate via radio with adjacent ATC
Providing radar-based control for sector	Total AERA failure : aircraft using backup clearances	NAS Stage A Procedures	Use PVD to plan Communicate via radio with aircraft and adjacent sectors

for assistance from time to time. (We shall discuss this possibility in more depth in Section V.) The ATC specialist, although managing the system, will nonetheless be outside the routine problem-solving loop. Will he be able to intervene, diagnose the situation, and in a timely manner solve the problem which the machine cannot? Will he be able to observe trouble spots the machine itself does not know exist through his routine monitoring? Research on human performance and our professional judgment suggest that this could be a very difficult task for a human [9,17,18].

If this role for the human seems problematic, consider the massive role alterations that will be instantly required if (when) AERA fails. According to the AERA Concept Document, a major portion of which is devoted to a detailed description of how AERA's fail-safe system would work, backup clearances guaranteed to be conflict-free for a short period would be continuously computed and stored. If a failure occurred, either because of actual AERA failure or an operator-initiated reconfiguration, these clearances would "drain" the airspace while other AERA failure-mode functions would reconfigure the center (or adjacent ones) for manual, present-day-style control.

Merely spotting such an emergency situation is difficult enough, even assuming the productivity gains expected by AERA designers are achieved (a factor of two or more). Reverting back to manual control may be impossible. AERA designers, recognizing this fact, intend for most of this backup function to be performed automatically, without human intervention. In the event of a massive AERA failure, the backup clearances would immediately become active, directing aircraft to contact adjacent centers for further control instructions, fly prescribed conflict-free (for 10 minutes, at least) courses out of the area of failure, or otherwise divert in the safest way possible. Flight plans would be sent automatically to the alternative centers, which would then assume control of the affected center's aircraft.

We do not know whether this fail-safe design will work; we perceive a high degree of uncertainty in it. It requires that either (1) the automated equipment will be able to handle virtually every aspect of the failure reconfiguration, or (2) the human operator will be able to reconfigure a system in which he is not actively involved. We think that many, if not most, failure conditions may be amenable to this plan, but such a combination of man and machine is extremely volatile and really not well understood today. Therefore, it cannot be described as a technologically conservative design approach.

So far, we have focused on only one failure mode, massive, large-scale, centerwide AERA failure. If state-of-the-art distributed computer architectures are used for AERA as planned, the probability of that event approaches zero. Much more likely is the failure of an individual

function (e.g., tactical execution or strategic planning). Different events may be considered "failures" in this context, each possessing intrinsically different levels of severity:

- A hardware device may fail. Most severe would be the entire suite of computers and their backups which perform one function; much less severe would be the failure of only one. In the worst case, the controller would be required to intervene until one of the redundant machines could be replaced or fixed; in the case of one machine failing, an automatic switch to an operative backup would occur, disrupting operations little if at all.
- A software system may fail or may be unable to handle a particular situation and will report that fact to the human operator. In this case, the operator might log the failure and intervene to resolve the problem. Presumably, this would occur routinely and frequently, since it is clearly impossible to program AERA to handle every contingency.
- A software system may function normally but perform its function inappropriately. In other words, it may fail but not know that it did so. This is the most insidious type of failure, and the type that will be the most difficult to detect and correct. Yet we would expect it to be the most common, especially during the early stages of AERA's existence. The controller will be expected to monitor and compensate for such deficiencies in AERA's programming.

Hardware failure is the easiest to deal with and prepare for, since backing up hardware with duplicate devices is relatively simple and inexpensive. Furthermore, detecting hardware failures is almost as easy as detecting massive centerwide failures; and accommodating them is either an automatic procedure or merely involves switching to the backup system(s).

Software deficiencies, planned or not, are more difficult to handle. Planned deficiencies increase the routine workload and training requirements of the controller, but they also perform a valuable function in that they require him to become regularly involved in what is otherwise a passive monitoring process. Unplanned software deficiencies can cause great problems, however, because they can disrupt operations at inopportune times. A good example would be the conflict resolver that reports a few seconds before a collision is about to occur, "Sorry. I've run out of memory and can't solve this one. Help!" Of course, some critical situations can be anticipated and planned for in advance (e.g., in this example, some lower-level separation assurance monitor or independent collision-avoidance system like ATARS could prevent the acci-

dent), but there will undoubtedly be other situations where software limitations surface at exactly the wrong moment.

In these cases, the controller will have to do the best he can to accommodate the failure. He might be able to pose the problem differently to the software, or he might see that it is not really a problem after all. But he may instead discover that in trying to handle the situation before reporting failure, the automatic problem-solver caused more problems than it solved. In any case, the controller must immediately make a transition from his monitoring role, intervene, diagnose the problem, and correct it in time. Whether it is intended or not, the controller force will be required to finish debugging AERA after it has become operational.

That debugging job becomes almost impossible when software failures of the third type occur. Not only must the controller back up a software system that can fail, he must watch that system to spot failures it cannot know about itself. The contradiction here, of course, is that even as the nominal traffic situation is getting so complex or large-scale that the human controller cannot handle it alone (per AERA productivity-improvement plans), the task of reliably monitoring the AERA control system as well is added to his responsibilities. Furthermore, even the most skilled controllers will be hard-pressed to notice these errors at all—and when they do, they may have no way of knowing what to do about them. The only way out of this trap is to create perfect software. We know that cannot be done.

SHARED CONTROL

Unlike AERA, the Shared Control system will be continually in transition. At every point, ATC specialists will be using a set of automated aids which may be coordinated manually or automatically. Over time, the number and capabilities of these aids will increase, until performance by about the year 2000 approximates or exceeds that of the AERA system. What distinguishes this scenario from the AERA one are (1) the means of arriving at this highly automated future, and (2) the degree of control continuously available to the human specialist as the scenario progresses.

Central to this scenario is the evolutionary introduction of increasingly powerful automated aids for the ATC specialist. This process has characterized ATC evolution so far, and we feel it should continue to do so. The problem of when to "throw the switch" to make the change to a fully automated control system is never encountered in this scenario. Instead, the human specialist gradually performs fewer and fewer mental and physical control functions as his automated assis-

tants become smarter and more numerous. He thus has a chance to ease gradually out of the role he knows and into that of subsystem activator and configurer. We have ended the scenario with even that function available to him automatically, but never is he denied the opportunity for active performance of *any* control function.

The ability of the human specialist to dynamically vary the allocation of tasks between himself and the computer is critically important to any future ATC system. In the Shared Control scenario, under routine conditions, specialists in the year 2000 will be able to assign some or all control functions for some or all aircraft to automated modules, retaining the remaining functions for themselves. They will be able to perform these functions much as they do today, or they will be able to selectively activate planning and monitoring modules as they desire. For example, in evaluating alternative trajectories for an aircraft, they will be able to use PAM to simulate future hypothetical situations, rather than relying on their own abilities to mentally simulate trajectories.

We must stress that these modules must be explicitly designed to be used in this fashion. If such capabilities exist only within a highly integrated automated package such as the AERA system, even the provision of sophisticated add-on man/machine interface packages may not enable their use in the fashion discussed here. Human needs must be given top priority during the initial design process.

The major rationale for allowing controllers to have flexibility in using the ATC computer system is that it enables them to maintain an optimal workload. In ATC and other complex control tasks, human performance degrades rapidly over time under either very low or very high workloads. Thus, in this scenario, in periods of low to very low airspace activity, specialists might perform many of the control functions for all aircraft or at least explicitly delegate functions to automation on a case-by-case basis. They will thereby remain involved enough to avoid lapses in attention that would impair their ability to recognize and respond to critical events in the airspace or in the operation of the ATC system. In periods of moderate air traffic, specialists will be able to assign more functions to automated control and perform only some functions themselves for selected aircraft. In heavy traffic periods, specialists will be able to provide required system throughput by assigning most routine planning and control to the machine. In this situation, their workload will consist of pilot requests that require their attention and highly selective intervention to modify or override trajectories planned by the automation system.

Although specialists could intervene in any control function, most observers believe they should rarely perform routine flight-track monitoring functions (e.g., delivering previously planned clearances at ap-

ropriate points, detecting deviations from expected trajectories) [9,12]. These functions are among the most time critical and must interface to independent conflict-alert or collision-avoidance systems. They will be developed and deployed earliest in the evolution of this comprehensive system of automated aids.

Planning functions, on the other hand, are less time critical and permit numerous opportunities for constructive human involvement. Any automated system will confront unusual airspace situations and system failures that require human intervention. ATC specialists must therefore maintain their skills for actively planning individual aircraft routes. They may decide to plan a trajectory from scratch or to modify a machine-suggested one. They may direct the various machine planning modules (the strategic or tactical planners, for example) to investigate and display the results of alternative "what if" options. They may opt to temporarily change the heuristics used for planning by selecting among built-in alternative strategies for plan generation. In short, they will be able to view the machine as an extension of their own planning expertise, instead of the other way around.

Although specialists will allocate primary planning or control functions, the machine will operate in a "shadow" mode at all times. A specialist planning and entering trajectories for aircraft may or may not use PAM for simulating future conditions and evaluating proposed plans. In either case, the plans he enters will be automatically evaluated for potential conflicts by the planning software, which will generate alerts when there are gross errors in planning. The human-entered plans might also be evaluated for efficiency relative to plans generated in the background by the automated planner, and advisories could be issued to the specialist when the planner determines that its solution is significantly better. Thus, the automated capabilities would serve as a redundant check on whatever functions the ATC specialist decides to perform, thereby reducing the potential effect of human errors and limitations. Unlike a system where the automation initiates solutions and requires the human to understand them and spot infrequent errors, Shared Control provides truly operational dissimilar redundancy.

The design serves a heuristic function as well. By actively generating solutions and then receiving performance feedback from the machine, specialists will achieve a better understanding of how the automated problem-solver works. This understanding will be invaluable when the system is operating in a mode where the automation is generating solutions and the humans are monitoring its performance, making them better able to perform the problem detection and correction functions described in the AERA context.

As a final consideration, we note that using the system in this way will enable specialists to maintain their control skills in the course of

normal activities. Requirements for expensive special skill-maintenance training will be significantly reduced, and we believe these gains will more than offset the occasional inefficiency in system operations that may be introduced when controllers take an active role in planning aircraft trajectories.

The main problem with this dynamic task allocation scheme is that designing such flexible automated components is very difficult. It requires the design of facile man/machine displays and input/output devices. Even more critical are the specific information requirements of the displays, keypads, and touchpanels. Extensive experiments must be performed to learn just what information the human specialist will need, how and when he will need it, and how to format and accept his actions in response to it.³ The challenge is formidable.

In summary, our Shared Control scenario provides several valuable features:

- Human-workload management to overcome the negative effects of workloads that are too low or too high on human performance and attitudes toward work.
- Dissimilar redundancy provided by automated checking of human actions.
- Opportunities for improved synergy between controller and computer, resulting in a better understanding of computer functioning by specialists.
- Skill maintenance through normal on-the-job activities rather than external training.

From the human-role perspective, at least, this scenario appears quite promising.

³See Appendix C for a preliminary set of experiment specifications.

V. TECHNICAL AND ECONOMIC IMPLICATIONS

In this section, we project technical performance and economic expectations for each scenario. By technical performance, we mean those aspects of each scenario which relate to the three primary goals of ATC: operational safety, fuel-use efficiency, and controller productivity. Where we can explicitly use quantitative information, we will carefully identify our information base; where we cannot, or where others have done so without sufficient justification, we will so state. By economic implications, we mean the cost and policy ramifications of choosing one pathway over another. Again, our projections are more qualitative than quantitative. Although our focus is primarily on the ultimate gains achievable under each scenario, we shall discuss interim systems where possible.

BASELINE

Extending present-day ATC practice into the next two decades will necessarily degrade overall system performance. All of the major goals of ATC will be increasingly compromised under the press of additional demand.

Consider the issue of safety. The largest contributors to system errors today are inadequate coordination between sectors, poor communication, and mistakes in judgment [19]. These problems will multiply as traffic increases. With denser traffic loadings will come smaller average sector sizes and more controller teams. Coordination requirements between adjacent sectors and teams will increase as the average sector transit times of aircraft decrease. Thus, a task which already typically consumes half of a controller's time¹ will consume even more. Communication channels, already overloaded in some areas, will incur even further delays and spur the use of improper terminology. With too many aircraft on a scope, a controller may not have the room or the time to ensure adequate separation.

Air operations will probably decline in efficiency as well. Although better three- and four-dimensional navigation systems are permitting point-to-point routings, the present archaic airspace structure and con-

¹Personal communication with Los Angeles Center and Los Angeles International Airport TRACON personnel.

troller cognitive limitations prevent their widespread use even today. Unless a means can be found to let the controller routinely plan and monitor irregular, point-to-point operations, ATC will continue to aggravate the fuel problem.

Efficiency of near-airport operations will also suffer under the Baseline scenario. Without building new facilities or providing wake vortex and wind shear advisory systems, decreasing inter-arrival spacings at the runway threshold will be the best way to increase capacity. That requires precise prediction of and adherence to strict threshold crossing times. And although the next generation of aircraft will possess flight management systems capable of such accuracy, without equivalent automated aids, controllers operating "by the seat of the pants" will not be able to use this capability. Again, a future ATC system similar to that of today will limit, rather than accommodate, the nation's air travel.

Finally, as mentioned earlier, controller productivity may actually decline under the press of smaller and denser sectors that require excessive coordination. In any case, we can safely assume that controller staffing requirements and manpower costs will parallel the rising demand for ATC services. By making some economic assumptions, we can arrive at a present discounted cost (PDC) for this scenario. Assume

- A linear growth in the controller staff from about 10,000 to about 20,000 by the year 2000.
- A 10 percent discount rate.
- A 7 percent per year salary increase.
- An overhead and fringe rate equal to 1.5 times salary.

Under these assumptions, the PDC of the Baseline scenario through 2000 is about \$4.5 billion. This figure considers only the increased cost of hiring more controllers. No allowance is made for other cost differentials—positive or negative—such as differentials in capital expenditures and fuel savings between this and the other scenarios.

However, these assumptions and predictions are countered by other emerging trends. Demand for ATC services may decline due to skyrocketing fuel costs and an inflation-induced drop in air carrier operations. Controller productivity has historically climbed over time as automated aids have become more sophisticated, and the advent of ETABS and similar equipment should further spur the individual controller's abilities to manage traffic with few automated decision aids. Projected increases in the controller work force may not materialize, and productivity may continue to increase. Figure 5.1 charts the climb in controller productivity from 1964 to 1979 [7] and extrapolates it to the year 2000.

If these latter predictions prove correct, our Baseline case PDC will

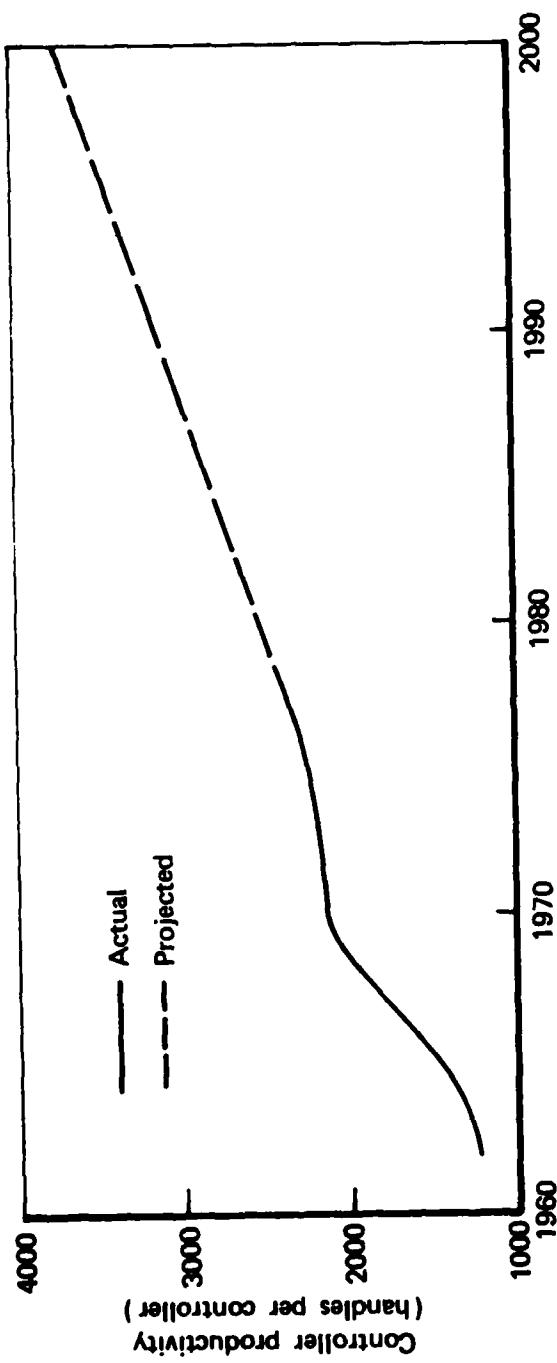


Fig. 5.1—Actual/projected controller productivity, 1960 to 2000

be considerably lower, resulting in a competitive cost position for it relative to the two alternative scenarios. While the Baseline scenario promises poorer performance, it avoids the significant R&D costs of the others.

Thus, support for the Baseline system will vary depending on the perspective from which it is viewed. Analyses whose main purpose is to show how well another scenario will fare compared to it will undoubtedly choose its worst manifestation, but advocates of the status quo might well use its better side to make their case.

AERA

Significant performance gains have been predicted for the AERA scenario. In Ref. 7, these predictions were quantified substantially. Two issues are addressed in this section:

- Can AERA's proposed automated capabilities really be achieved within the time frame considered here?
- If so, will the predicted performance gains and cost reductions result?

We will address primarily the "ultimate AERA" system, since this scenario discounts interim performance gains. According to E. Koenke of the FAA's Office of Systems Engineering Management [20]:

AERA will be a fully automated system, embedded within each en route facility, that will (a) automatically plan conflict-free, fuel efficient profiles for aircraft operation in positively controlled airspace, (b) generate ATC messages needed to execute the planned profile and assure aircraft separation, and (c) deliver ATC messages via a data link or VHF voice channel.

What assurances are there that AERA is technologically achievable? In engineering development tasks like this, standard procedure is to construct a feasibility-demonstration system first to show that the especially hard subproblems can be handled in the lab. This was indeed done [5], but the resulting system neither demonstrated feasibility of the problem-solving algorithms that were implemented nor confronted the especially difficult tasks. The AERA feasibility-demonstration system

- Utilized an airspace simulation in which the aircraft maneuvered precisely as ordered by the automated controller.
- Did not employ the more difficult resolution techniques of hori-

zontal vectoring or alternative routings, but instead used fixed routes and timed descent points, altitude, and speed.

- Used a pairwise conflict-resolution algorithm untested in complex situations involving many potential solution choices and variable pilot responses to clearances.
- Did not interact with human controllers or pilots at all.

Although this work did implement some basic AERA tasks using a real sector environment and data, it was based on highly simplified problem-solving techniques.

One might argue that although the existing feasibility demonstration is not convincing, the problems facing AERA designers will surely succumb to the tests performed when the testbed is completed. Perhaps. It appears to us, however, that some very basic research results are required before AERA can be responsibly supported as a feasible next-generation ATC system whose engineering development can be started today.

Take, for instance, the simple task of handling routine operations.

Example 1:

"Let's take all the northbounds and vector them via J13/J25 to avoid that thunderstorm cell lying in the middle of their usual J20 route."

This sort of operation typifies day-to-day controller decisionmaking. The situation requires that a response be applied to a group of relevant aircraft. Present-day controllers notice such things routinely and quickly adapt their actions accordingly. Yet this case could be quite difficult for an AERA system to handle, unless the system were specifically preprogrammed for it. Through either controller directives or its own weather-avoidance subsystem, AERA would have to alter the parameters of strategic and tactical planning to avoid the severe weather. Consider how many real-world concepts AERA must "know about" in this case:

- *Northbounds.* Aircraft may be classified in various ways, including by general direction of motion. Generalized, the notion of "class" would require many pattern-matching capabilities to extract subsets of active aircraft based on numerous characteristics: geographic area, similarity or dissimilarity of capability, proximity to one another, flight status, and so on. Once a class has been defined, AERA must then be able to apply a set of operations to each member of the class, rather than to a single aircraft only.

- *Via J13/J25 vs. normal J20 route.* AERA must know how to use alternative routings to achieve the same destination or a nearby destination (how near is near enough?), without too badly violating in-place constraints for the flight or the route (how much violation is too much?).
- *Thunderstorm cell.* Of course, AERA will "know" what severe weather is and how to avoid it, but will it be able to switch high-level strategic plans according to developing weather patterns as demonstrated here? And if it can do this automatically, how consistent will the planning algorithms be? Will oscillations occur as the automated planner teeters at the edge of a planning strategy, switching back and forth according to minute changes in the environment? "Focus of control" and "consistency of strategy" are well-known concepts in artificial-intelligence planning research, but they have never been applied to a real-life planning system.

As currently designed, AERA would possess none of the above concepts. Without the concept of class, AERA will alter each flight's profile individually, resulting in a haphazard but "optimal" (from the machine's perspective) set of unique profiles around the weather. Without alternate routes, AERA will issue potentially complex vectors around the weather instead of anticipating its movement and adapting traffic flows in general. Without the concept of high-level problem-solving strategies (an embryonic research subject even in the field of artificial intelligence), AERA, like other problem-solvers of its type, may display highly erratic behavior that is impossible for a human to understand.

This example is quite simple compared with what AERA could be faced with. Consider another:

Example 2:

"Hey! What's that military aircraft doing? It's flying erratically at supersonic speeds. Scatter the traffic! Twenty mile buffer."

Admittedly, this is a low-probability situation, but it may be more likely than the massive centerwide failure discussed in Ref. 7. Despite its improbability, this example does illustrate a very important capability that a fully automated ATC system must possess. While this situation could be readily detected and handled by a human controller today, under full AERA, not only would detection be questionable (because of

human monitoring limitations), but handling the situation while remaining "out of the loop" could be extremely difficult. Will AERA be programmed to spot such out-of-bounds situations and handle them, or at least to alert the system manager? The performance of other comparably complex automated systems, such as nuclear reactor control systems, testifies to the fact that many potential emergency situations will fail to be taken into account by system designers.

One way out of this trap is to design the software system as a set of independent modules which can be used selectively by the human operator. Then, if some of the modules are ineffective, the operator still has aids available to facilitate his handling of the unusual situation. However, AERA is not currently conceived for use in such a modular fashion. To be used independently, each module must have a well-engineered human interface, and AERA, by emphasizing completeness of capability, will attempt to circumvent this requirement with a monolithic full-capability system. Its human interfaces will allow manipulation of the system as a whole, but not as a set of independent modules.

Even if AERA could work as claimed, this emphasis on functional completeness will necessarily reduce its impact. Virtually all aspects of the current controller's job must be handled by AERA at some level of expertise for it to operate as autonomously as designed. This in turn requires the selection of an environment that permits universal application of machine problem-solving skills, an environment simple enough to handle automatically. For AERA, this environment must be the en route high and transition sectors, where only well-equipped aircraft fly, surveillance lapses are infrequent, and everyone is on flight plans. In short, control in these sectors is routine and simple compared to the complexities inherent in the rest of ATC.

But en route high and transition sectors are not where the problem is. While some gains in fuel efficiency can probably be achieved there, AERA's anticipated controller productivity gain of 100 percent would net savings of less than 8 percent in overall ATC personnel costs because of the relatively few controller positions that would actually be vacated under AERA.² ETABS' 25 to 30 percent productivity increase for all en route sectors³ would overshadow these savings at a far lower cost, with almost no technological uncertainty. Furthermore, terminal control and en route low-altitude sectors will become overloaded long before an AERA-like system could be applied there. Even if AERA becomes standard equipment in its applicable arena, increases in controller staffing to accommodate other shortfalls will more than offset its gains.

²Personal communication with Glenn Kinney, The MITRE Corporation.

³"Electronic Tabular Digital Display (ETABS)," briefing by J. Edgebert of the FAA, presented to The Rand Corporation in February 1980.

Likewise, AERA may net little in the way of fuel savings. AERA advocates point to a typical New York/Washington flight profile in which procedural altitude limitations cause inefficient fuel usage. In the past, these procedural restrictions caused a 7 to 8 percent fuel penalty, but recent relaxation of them has reduced that penalty to 3 percent in heavy traffic periods. It appears that even in this extreme case, a 50 percent or greater improvement in fuel efficiency could be achieved without AERA. As to the remaining 3 percent inefficiency, any conflict-prediction aid should improve this situation, leaving unanswered the question of just where a full AERA system would uniquely contribute to fuel savings.

In conclusion, AERA may indeed improve system safety if it is fully and correctly implemented. Human lapses of judgment, coordination, communication, and attention may disappear when a fully automated AERA system is unveiled. However, these uniquely human problems may be replaced with uniquely machine problems such as regimented responses to novel situations, uncontrolled "free running" operation due to decay of operator monitoring skills, or increased rather than decreased continuing costs.

SHARED CONTROL

The Shared Control scenario appears to trade off some of AERA's uncertainties for lowered performance expectations. Each of the aids postulated for this scenario seems technologically feasible for the deployment time frame suggested. The main questions we must address concern the specific capabilities and human interfaces provided by the aids and the human specialist's abilities to capitalize on them.

For example, in the Shared Control system a digital communications regime is implemented with a Tactical Communications Manager (TCM) to handle it. We know that automatically monitoring and controlling the execution of stored plans will dramatically reduce a component of the ATC specialist's workload that depends directly on traffic density. However, a series of pilot experiments performed at Rand using a highly simplified ATC simulation indicates that by itself a TCM may not provide any great improvement in a controller's ability to handle increased traffic loads. We provided a rudimentary TCM to half the subjects in our experiments and found that objective performance differences such as total fuel use or number of conflicts yielded no significant differences between groups with and without TCM. Those using a TCM were able to handle low-density airspaces with reduced

workload, but they resorted to strict manual control in moderate and high-density situations. At high densities, the unaided controller simply cannot plan conflict-free clearances for storage in the TCM.

There is some evidence, however, that providing a planning aid in conjunction with a TCM will improve an ATC specialist's ability to handle more traffic. Preliminary findings from Great Britain, which has tested an aid similar to our strategic planner, indicate that this technique does indeed appear to "level out" varied traffic loads and facilitate the handling of more traffic [15].

But what does "more traffic" mean—denser traffic loadings within a certain fixed-size sector, or larger sectors with no density changes? The answer depends on the distribution of the specialist's workload and can only be answered experimentally. Our limited experience indicates that a TCM will mainly facilitate the management of larger sectors rather than denser ones, because a TCM reduces the specialist's monitoring load rather than his planning load. Other aids, such as a strategic planner or Executive, would moderate his planning workload and allow denser or more complex sector traffic configurations.

This example illustrates that ATC aiding/automation components may selectively enhance different features of system performance in specific ranges of system load. Likewise, we would expect differential effects for different types of airspaces. A component's effect clearly depends on other available components and on whether they are used as aids or as stand-alone automation. Appendix C outlines a program of empirical research into the effects of various combinations of aids for the future ATC specialist.

We must also point out that the Shared Control system may not achieve the peak productivity gains of a perfected, fully automated system like the ultimate AERA for routine operations. Keeping a man in the loop entails some cost. It means that he must continuously comprehend the developing traffic situation well enough to react quickly and appropriately. Having automated planning and monitoring tools will greatly reduce his cognitive load and increase the average number of aircraft he can oversee, but his mental capacities will at all times govern how far this increase can go.

However, the flexibility afforded by allowing the specialist to allocate tasks between himself and his supporting technology should produce a system that can handle unusual cases efficiently and accommodate the desires of individual traffic under light to moderate overall system loads. As mentioned in Section IV, this scenario should enable controllers to maintain their skills at the high level needed to meet requirements for controller proficiency in failure-mode operation. Regular specialist experience with different aiding/automation configurations should especially reduce the perturbations in ATC oper-

ations caused by failures or problems that arise when the new technology is applied to specific situations. Investing in interfaces to support each component's regular use in different configurations should also contribute to stability of operations during partial failures when specialists must assume certain functions.

We would also expect the more moderate performance gains promised by the Shared Control system to be balanced by its wider applicability. Each new system component should be readily adaptable to terminal control and mixed airspaces in addition to positive-control en route airspace. Each component should be deployable by itself whenever it reaches a suitable stage of development without having to wait for a complete system of modules to be proven fully competent. More than anything else, this aspect of the Shared Control scenario—its far-reaching applicability—places its overall benefit/cost ratio far above that of AERA.

Overall staffing costs in this scenario should be significantly below those in the other two scenarios because of the incremental deployment of automated aids. Figure 5.2 charts our staffing predictions for the three scenarios. We anticipate that staffing levels will climb roughly in step with traffic growth in the Baseline case and the AERA case, with the AERA system realizing a substantial personnel reduction in the en route centers when it is finally fielded. Given its limited applicability and the distribution of controllers between terminal and en route control centers, its year-2000 staffing requirement should fall somewhere below that of the Baseline case and above that of Shared Control. After some start-up costs, the steady introduction of automated aids in the Shared Control system is assumed to result in a relatively constant level of specialist staffing in both the en route and terminal-area centers.

Decreased fuel use should also produce cost savings in the Shared Control system. Like AERA, Shared Control will relax procedural constraints and allow more optimal flight profiles. Due to our requirement for keeping the human controller in the control loop, however, we anticipate that many restrictions would be retained for his benefit, thus reducing the possible cost savings. Whether Shared Control or any other system would achieve a fuel savings of 0.1 percent or 10 percent cannot be known without a careful analysis of which current procedural restrictions could be relaxed under each postulated system.

The development and implementation costs in this scenario are similarly elusive. Whereas the Baseline scenario involved none of these costs, Shared Control will require a substantial investment in both basic research and engineering development, given the sophistication of the planned aids. We cannot anticipate just what these costs will be, but they should approximate those for AERA development. (Although

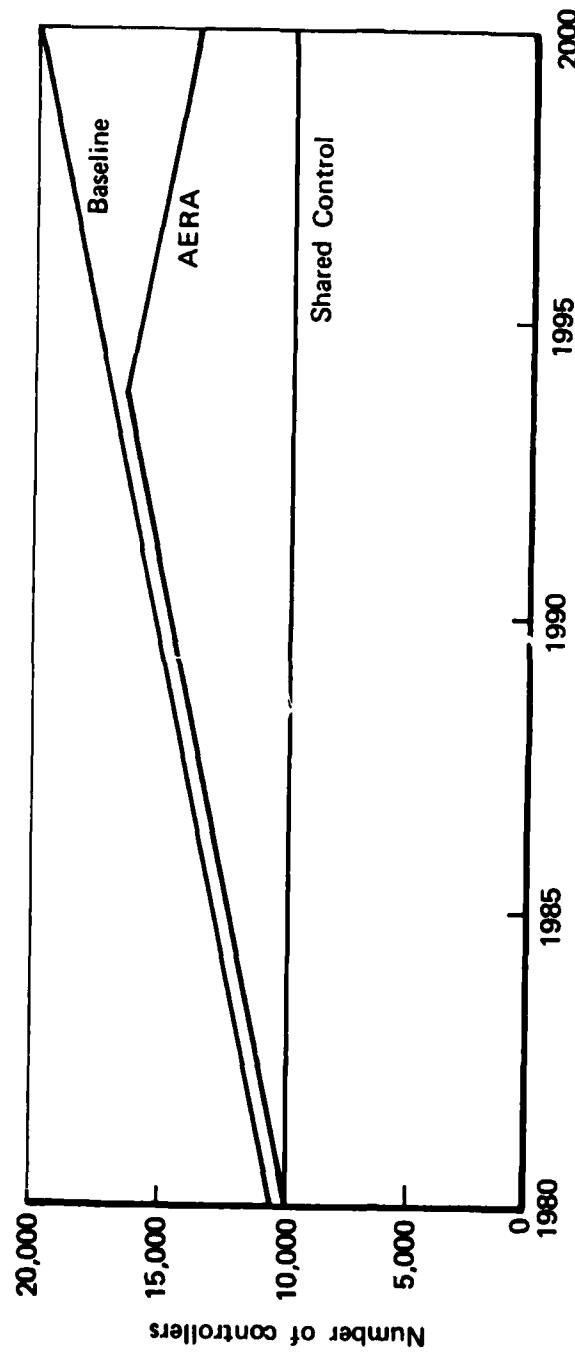


Fig. 5.2—Anticipated staffing levels under alternative scenarios

less software appears necessary to implement the functions, more effort will be required to design the more complex human and machine interfaces. Hardware requirements are comparable; however, this scenario should entail a lower total manpower cost than the other two scenarios over the next 20 years. Because AERA would be deployed as a complete system, the controller force would have multiplied significantly by the time the system was actually fielded. Thus, overall interim costs of Shared Control should fall below those of the AERA scenario, even though the development costs are comparable.

In summary, Shared Control will provide incremental performance enhancements for different traffic situations in a range of airspaces throughout its development. It may trade off maximum productivity gain (relative to a more highly automated system) in normal en route operations for improved flexibility in responding to failures and special situations. Thus overall productivity gain or cost savings for ATC services and users of those services may well approach those projected for AERA while providing a more robust system that fully uses human and computer capabilities for mutually redundant backup.

SUMMARY OF THE TECHNOLOGICAL ISSUES

From a performance perspective, simply adding more controllers to a present-day system is ultimately counterproductive. In contrast, AERA attempts to achieve the maximum performance possible by relying on highly advanced but highly uncertain technology. Because of its emphasis on functional completeness to remove the human specialist from the routine control loop, it must compromise in its domain of applicability. And because of this compromise, its ultimate effect is significantly depressed. The Shared Control scenario, on the other hand, attempts to create a symbiosis of man and machine in which each is responsible for a carefully defined subset of the ATC task, with the human dynamically determining the exact distribution of these responsibilities. While the aids suggested for a Shared Control system may not prove to increase performance in en route control as much as does the AERA system, they are designed with a wider range of domain applicability and flexibility. Thus, their overall effect is likely to be greater than AERA's, and there is considerably less uncertainty that they can be developed.

The AERA and Shared Control systems both require state-of-the-art technology; and each requires a better understanding of a common set of human/machine system design issues. In particular, four general

categories of technical questions are critical in determining how to best evolve either system or a combination of both:⁴

1. *Algorithmic choices.* What are the best forms of automated information-gathering, data representation, and problem-solving? Are the proposed planning structures stable in all situations? How does interfacing with CDTI affect the planning algorithms?
2. *Control/display requirements.* What operator inputs are required and what information should be displayed? What communication protocols must be defined for voice and digital transmissions?
3. *Demands on the human operator.* What functions are required of the human operator in the normal and failure modes of system operation? What minimum performance levels are required? What workloads are estimated?
4. *Failure backup.* What types of human and machine response are required to deal with each possible subsystem failure? What are the capabilities of the human in system monitoring and offloading? How much involvement in normal-mode operations is necessary to maintain human intervention proficiency?

In this section, we shall review each of these issues individually. Appendix C presents a suggested experimental program which may begin to resolve them.

Algorithmic Choices

The AERA and Shared Control scenarios are based on a myriad of planning and control algorithms, some now in operation, some in development, some to be produced years from now. Many techniques exist for designing such software. Different forms of route planning and replanning may be favored in different situations: high- or low-density airspaces, open or restricted communications, long or short time-response requirements, etc. Planning options include

- *Individual sequential planning.* This is the simplest form of planning. When an aircraft is about to enter the airspace, its flight plan is converted into a precise four-dimensional profile. If the flight plan conflicts with that of another aircraft, pre-stored deterministic rules of conflict resolution are used to alter

⁴This section is primarily addressed to the developer of an automated system and thus is presented in somewhat more technical language than the rest of the report. The reader may wish to review the AERA Concept Document [7] as background to this material.

it appropriately. The flight plans of other aircraft already in the airspace are typically not modified unless absolutely necessary. The initial AERA planner should function in this fashion.

- *Multiple-aircraft coordinated planning.* Instead of planning for only the entering aircraft, the program may modify the flight path of any aircraft.⁵ It first checks to determine if an efficient route is possible without movement of other aircraft. If not, the profiles of existing aircraft are altered until an efficient, optimal, conflict-free set of paths is determined. The solution that results should minimize fuel and time for all aircraft rather than penalizing aircraft that arrive late at a congested airspace. Of course, this technique may require more processing time than the simpler individual form of planning and may generate an unacceptable rate of clearance revisions.
- *Sectorwide replanning.* In the event of a severe weather disturbance or an airport closing, virtually all aircraft within the sector may have to be rerouted. This may require libraries of backup plans for each possible circumstance, which can then be adapted as necessary. Also, high-level multiple-aircraft commands may have to be defined. This planning technique has been investigated and reported in Ref. 21.

Another issue of algorithmic choice concerns the various functional hierarchies within any automated ATC regime—whether the system functions as a set of individual modules or as a fully integrated single system. There will always be a number of interacting subfunctions such as flow control, metering, rerouting, strategic planning, and conflict avoidance. Weather fronts, runway closings, emergency vehicles, or controller or pilot intervention may destabilize these interactions. The experimental program must consider the multiple subsystems functioning together as well as individually.

Control/Display Requirements

Once the general approach to planning and control algorithms is determined for each interim system, the necessary controls and displays can be designed. The major design issues include:

1. Determining the necessary operator inputs. These may include:
 - Modifications to potential or actual clearances.

⁵"Electronic Tabular Digital Display (ETABS)," briefing by J. Edgebert of the FAA, presented to The Rand Corporation in February 1980.

- Modifications to airspace restrictions.
- Requests for explanation.
- Requests for diagnostic information.
- Requests to transmit data.
- Requests to reconfigure displays.
- Declaration of open-loop operation.
- Allocation of control.
- Changes to system evaluation functions.

2. Determining the types of information to be displayed. Possible displays include

- Static planning displays giving path snapshots with translate and zoom capabilities. The operator may also need capabilities for decluttering through selection of subject and object aircraft.
- Graphic planning displays giving dynamic path trajectories. The operator may have control over rate.
- Graphic planning displays highlighting aircraft densities over time (needed for metering/flow control).
- Graphic planning displays highlighting resectorization and indicating boundary changes, established plans of new aircraft, traffic advisories, etc.
- Tabular data displays giving command prompts for verbal communication to pilot.
- Tabular data displays alerting the controller to system errors or other requests for intervention. (This may be coordinated with additional information presented on a graphic planning display.) The information to be presented includes coasting clearances, degree of system stabilization, and predicted downtime. Also, in normal operation, the system should show its immediate and projected loading levels.
- Tabular data displays showing direct pilot-to-system interactions. Special protocols may need to be defined for opening such communications and for informing the controller of any plan changes that result.

Demands on the Human Operator

An extremely important input to the design process is the loading on the human operator. This information is difficult to ascertain without access to a pre-production system for task analyses. In the absence of this, task analyses will have to be performed on the FAA's current

testbed, newly constructed abstracted simulations, and conventional ATC systems.

The information that should be collected includes

- The proportion of traffic handled by each of the forms of active planning—strategic planning, conflict-avoidance patching, flow control, weather avoidance, out-of-association replanning, emergency response, and pilot request replanning.
- The distribution of controller planning behavior between open- and closed-loop control. This work should produce a table of individual activities, their timing, their demands on the human controller, and the branching criteria to other activities.
- The amounts of controller monitoring devoted to separation assurance, track association, weather conditions, flow control, system status, and intersector coordination. We also should estimate transition losses incurred as the controller switches among these tasks.

In general, a metric is needed for rating the various operator tasks in terms of their load level. We must know more about the relationship between load levels and operator performance (time delays, error rates). In the end, we need to estimate required staffing levels and performance of the human component of each candidate system configuration.

Failure Backup

We know very little about the capabilities of the human operator to perform failure backup. Some of the displays necessary for human backup were described in the control/display section. Here, we concentrate on the testing of procedures for different types of system failure. These conditions include

1. *Full-system failure.* In either of the highly automated scenarios, the specialist will become more active in the control process when a massive failure occurs. In a full AERA system, he will participate in the reconfiguration process and may even be required to manually control traffic (his role has not yet been precisely determined for this scenario). In the Shared Control case, whenever an automated-aid failure occurs, the specialist is expected to step in and perform that aid's function himself. In each of these cases, the following characteristics of this role change should be studied:
 - The operator's ability to assess the current and projected state of the airspace while monitoring. This must be evalu-

ated for different airspace sizes, aircraft densities, and aircraft mixtures.

- The amount of operator activity (active aircraft control, system queries, etc.) necessary to maintain controller proficiency.
- The time lags involved in shifting from a passive to an active controller mode.
- The relative importance of the following procedures for reducing the operator load during a system outage: plan coasting, increased protection areas around the aircraft, emergency flow control, and resectorization.

2. *Partial-system failure.* Failures may occur in the separation assurance, planning, metering, flow control, execution, monitoring, data link, or any other system modules. Some of these failures will be more severe than others. The interaction among module failures and human response is still poorly understood.
3. *Offloading.* Each system will occasionally request the operator to take control of some subset of the aircraft in the airspace. We need to determine the capacity of the operator to take over control of such aircraft and simultaneously monitor the critical system functions.

VI. CONCLUSIONS

We have considered several alternative ATC futures, beginning with a Baseline case in which nothing beyond the most conservative R&D projects paid off. We have concluded that the approach of simply adding more and more controllers is ultimately counterproductive from a performance standpoint. We have examined the FAA's plan to use advanced computer science technology to construct a fully automated ATC system for application near the year 2000. The expected aircraft safety levels, fuel-use efficiency, and controller productivity have led us to question that plan and to suggest that there may be a middle ground consisting of a highly, but not totally, automated system.

We believe that pursuing the goal of full-automation AERA—with little regard for interim systems or evolutionary development—is a very questionable R&D strategy for ATC. It seems unlikely that a large-scale multi-level AERA system that can effectively handle non-routine events, show stable behavior under dynamically changing conditions, and be virtually immune to reliability problems can be implemented in the foreseeable future. Human controllers may be required to assume control in at least some of these situations, although at present there is no conclusive evidence that they would be able to do so; indeed, some evidence and opinions from the human-factors community suggest that they would not be able to.

The AERA scenario presents serious problems for each of the three major goals of ATC—safety, efficiency, and increased productivity. By depending on an autonomous, complex, fail-safe system to compensate for keeping the human controller out of the routine decisionmaking loop, the AERA scenario jeopardizes the goal of safety. Ironically, the better AERA works, the more complacent its human managers may become, the less often they may question its actions, and the more likely the system is to fail without their knowledge. We have argued that not only is AERA's complex, costly, fail-safe system questionable from a technical perspective, it is also unnecessary in other, more moderate ATC system designs.

Some AERA advocates assert that it is necessary to keep the human out of the time-critical loop to achieve productivity and fuel-use gains. We question that belief as well. AERA may well achieve 100 percent productivity increases in the en route high and transition sectors, and it may indeed facilitate more fuel-efficient air operations. But if the controller work force almost doubles, as expected, by the time AERA comes on-line, and AERA's domain of applicability is limited to the

simplest of sector types, its ultimate effect may hardly be felt, since the actual ATC bottlenecks occur elsewhere. Further, greater fuel efficiency comes from many sources—some as simple as present-day relaxation of procedural restrictions, some as complex as the planning modules of AERA and Shared Control. AERA may meet the goals of ATC by 2000, but the costs incurred along the way will be very great—in dollars, in fundamental research that must be completed, and in restrictions on the controller's role.

Ultimately, the AERA scenario troubles us because it allows for few errors or missteps. The right choices have to be made at the right times, or a failed AERA scenario would degrade to a more costly and delayed version of the Baseline scenario. In the attempt to construct a totally automated ATC control system, unacceptably high possibilities and costs of failure overshadow the potential rewards of success.

Our main conclusion is that such an overwhelming dependence on technology is simply unnecessary. If the planned AERA scenario were altered only slightly, it would be essentially equivalent to the Shared Control scenario. All of its technical building blocks are present in Shared Control:

- Air/ground datalink communication.
- Strategic planning (profile generation and alteration) and operator displays.
- Tactical execution.
- Track monitoring and alert.

Missing, however, is the right principle for piecing these building blocks together. Under AERA, they would be fully integrated into a single problem-solving system which extends its capabilities by infrequently requesting human action; under Shared Control, the building blocks would themselves be extensions of human capabilities. Operationally, this shift in perspective requires two modifications of AERA plans:

- The role of man under AERA would be expanded so that he is routinely involved in the minute-to-minute operation of the system.
- The system would be constructed as a series of independently operable, serially deployable aiding modules.

The state of the art in ATC problem-solving techniques does not validate the minimal AERA human role; neither does established knowledge about human limitations or capabilities in this domain. Insisting that man be essentially automated out of such a critical control system is an unnecessarily high-risk approach.

If the system is designed to support him, we would expect the future ATC specialist to take a very active and creative role in manipulating his aiding modules. Safety could be assured by assigning the machine primary responsibility for routine separation assurance tasks at the lowest levels. The specialist should be responsible for comprehending situations at high levels of abstraction and activating modules to meet the ever-changing demands of those situations. He should be able to adjust a module's parameters and its relationships to other modules so that instead of simply monitoring the machine's preprogrammed sequence of instructions, he actually controls the outcome. He should be given the authority to determine which operation the machine performs and which he performs. He should be given the opportunity to learn all of this gradually and to influence the system's design before it is finalized.

This shift in perspective captures the spirit of this report. Specifications of module capabilities and their sequence of implementation are best left to designers who are intimately familiar with the engineering details. We have presented just one of many alternatives in which man has a significant ATC role; the details of the system design need refinement and may indeed undergo great change in the process. For example, our Shared Control scenario suggests implementing digital communications before providing any planning aids at all. Perhaps events will dictate otherwise—a late DABS introduction and an early development of automated planning techniques could reverse this sequence. Fielding a planning aid first as a stand-alone module would not compromise the Shared Control scenario in any way. The essence of the Shared Control scenario is reflected in its name—man and machine must work together and *share* in the overall control function of ATC.

Our key concern is that the human specialist's unique capabilities be acknowledged and the technical uncertainties of an AERA-like system be recognized and dealt with before too much of the Baseline scenario comes to pass. If this is not done, we risk relying solely on an unproven, costly technology to meet the nation's demands for ATC service. We have shown not only that there is a feasible alternative, but also that this alternative may result in lower costs, a higher level of performance, and a more satisfying role for the personnel who will be responsible for moving air traffic safely and smoothly.

Appendix A

SATELLITE-BASED AIR TRAFFIC CONTROL

This appendix presents a high-technology scenario based on several ATC systems recently proposed by the aerospace industry [13]. It envisions a future of nationally centralized satellite-based ATC. It is a "throw-away-the-book-and-design-from-scratch," ultimate ATC system based on satellite technology. This system would

- Allow an equipped aircraft to pinpoint its position anywhere over the continental United States more precisely than can be done using the current VOR navigational system.
- Allow an equipped aircraft to communicate with one of several regional control centers, regardless of its proximity to them.
- Provide datalink as well as voice communication.
- Be more reliable than the current distributed system of ground-based navigational beacons and ATC centers/sectors/terminals.

Concurrent with the development of the satellite system and its attendant user modules¹ will be the design and construction of two or more regional control centers (RCCs). Perhaps in conjunction with a continental control center (CCC) to back them up and provide centralized flow control, they will replace all existing en route control centers and many terminal-area approach controls as well. However, most terminal radar control areas will remain, although their surveillance data input may be replaced with satellite-derived and RCC-forwarded information.

At least three classes of airspace users will be defined under this system. Lowest on the totem pole are "uncontrolled" aircraft, with no equipment requirements. They may not enter controlled airspace, which generally exists around any terminal area of any size at all and everywhere above 3000 feet. Next comes a class of users called "cooperative." Most present-day VFR and minimal IFR users will be in this class and will be required to have a minimal complement of satellite-navigation and communication equipment. They will be permitted to fly in controlled airspace during off-peak hours and will be afforded

¹The airborne navigation and communication equipment and the ground-based air/ground and ground/ground networking equipment.

separation assurance and navigational services from ATC. The third class, "controlled" aircraft, must be equipped with highly capable satellite-navigational processors and complete backup systems. For this considerable cost outlay, they will be given preferential treatment in controlled airspace at all times and will be guaranteed arrival slots at the major terminals.

The initial designs call for extensive automation of the control process in a manner similar to the AERA scenario. Sectorization will divide the national airspace into pieces small enough to be managed by the many human/computer teams at each center, as before; however, since this is a "from-scratch" design, pre-existing airspace limitations, shelves, and routes from the ground-referenced days of VOR will be ignored, and new ones will be created where necessary. Flexibility in point-to-point operations is the watchword. Also stressed are the backup procedures which transfer control from center to center, or RCC to CCC, in the event of massive failure. Minor, function-specific failure is less of a problem, for multiple-processor hardware designs are employed.

Why have we dismissed this scenario out of hand? To begin with, it is clearly the most revolutionary of all those considered in this study. It requires changing virtually every aspect of the air traffic service, from basic navigational aids and air routes to the basic concept of a localized, distributed ATC system. It implies replacing existing ATC centers with two centralized RCCs; replacing the present surveillance system with a satellite-based one; replacing VHF with C- and L-band satellite-based network communications; replacing VOR-based navigation with a satellite-based system. When complete, every ATC-related box will be replaced in the air and on the ground. While this replacement can be accomplished gradually, the endpoint is a completely new ATC system.

That might not be so much of a problem if the new ATC system were clearly feasible and obviously better than any of the other alternatives. It probably could be built. It might be a very good system. But the uncertainty on both counts is so high that betting on it is betting against the odds. Such a system would require at least

- A satellite-based navigational system which lies just at the state of the art today. In fact, the problems, both technical and political, that have prevented the widespread use of GPS-based navigation would also apply here.
- A satellite-based communications network rivaling the nation's telephone system. To our knowledge, the complexity and reliability requirements of such a network surpass those of any known analogous system.

- A computer-based ATC problem-solving system with all the capabilities envisioned for AERA, but an order of magnitude larger.
- Incentives or rule-making actions to lead private industry to build and aircraft owners to buy significantly more expensive and technically sophisticated cockpit equipment.
- A backup plan or system an order of magnitude better than that required for any of the other scenarios, since the failure of an RCC would affect at least that many more aircraft.

Appendix B

ELECTRONIC FLIGHT RULE AIR TRAFFIC CONTROL

Our final scenario departs from the others by applying advanced technology primarily in the cockpit. Taking recent FAA-sponsored research on this topic as a starting point [22] and using current research in distributed artificial intelligence as a guide [23], we have constructed a scenario to implement the assertion that control left the cockpit for technological reasons (the availability of ground-based radar to "see" through weather) and can now be returned via similar advances ("intelligent" CDTI and CAS). In other words, the overriding philosophy of this scenario is to move as much of the control process as possible—including separation assurance and flow control—back to the individual aircraft and reduce the ground controller's involvement to a minimum.

How would this philosophy be operationalized? What does "get control back into the cockpit" really mean? In some instances, it would mean that an aircraft could fly to its destination in IMC without even having to file a flight plan. This would be the norm in uncongested areas. Most flights, however, would involve some coordination between air and ground and thus would require flight plans and ongoing air/ground interactions. The difference between the ATC system of this scenario and that of the Baseline case lies in the degree and kind of ground involvement in an aircraft's operation. Where our Baseline case would continue current ATC practices of assigning altitudes, headings, and perhaps other parameters of flight, this system would have aircraft automatically, digitally "negotiating away" conflicts, using cockpit-located black boxes. This would occur with a minimum of pilot or controller intervention.

In this scenario, emphasis is placed on methods available for gradually converting the current ATC system into a primarily electronic flight rule (EFR) system: Controller/pilot duties are redefined, protocols devised, hardware and software requirements designed. A network of ground-based transmission facilities necessary for CDTI systems is specified. A program is begun to create the requirements for a uniform, approved, intelligent CDTI/CAS system. Initially, such a system will emphasize known conflict recognition and resolution technology, but a long-term R&D program will be started to design the algorithms for cooperative separation assurance and flow control. Perhaps the most

important program will expand a current NASA research program [11] aimed at discovering just how such a hybrid ground/air-based control system would function—whether it could handle, in a highly decentralized way, general problems such as flow control. Extensive testing using laboratory simulations will be performed once the initial design is stabilized. Preliminary experiments of this nature indicate that the limitations of EFR will stem primarily from

- Workload limits during single-pilot operation.
- Uncertainty near ground-control sector boundaries.
- The inability of EFR problem-solving, whether human-assisted or not, to produce efficient flow-control solutions.

The third problem requires a redefinition of the role of ground control to that of high-level arbiter and flow controller; the first implies that single-pilot operation be conducted under traditional IFR procedures or only in uncongested airspace; the second will remain to plague implementers for quite a while.

Under this scenario, controllers continue to perform many of their current-day tasks, but their normal operations are skewed toward general, abstract flow planning, tie-breaking, and decisionmaking under unusual conditions. The expansion of EFR into even congested areas may cause more problems than anticipated with inefficiency and undesirability of machine-generated solutions. This in turn may cause pilots to request the intervention of controllers more frequently. In order to properly handle such situations, controllers must, of course, be cognizant of these encounters, so their skills at rapidly "coming up to speed" on a situation they were not routinely monitoring must be finely honed.

The controllers' automated assistance must change with their evolving role. Automatic target acquisition, flight-plan query, ground/air communication, monitoring, and prediction will enable the controller to know the intentions of the EFR aircraft in his sector and to intervene easily if necessary. The controller will service each aircraft crossing his sector according to its equipage. To the poorly equipped or non-equipped aircraft, he will provide traditional IFR services (although rumors continue that such services will be discontinued in the near future); he will insure that EFR-equipped aircraft know about all aircraft in his sector, and he will intervene as discussed above.

Is such a role possible? What assurances are there that such a mixture of control responsibility would work? Answers are just now appearing in the literature. Joint FAA/NASA work [11] has shown that controllers can learn to handle simultaneous EFR and IFR operations, at least under simplified laboratory conditions. Although the effects of such a system on total workloads are uncertain, the controllers reported

having to work harder on the IFR traffic they did control because of the uncertainties about exactly what the EFR aircraft were doing. We would expect this increase in workload to generally be more than offset by the reduction in total targets being controlled, resulting in generally reduced workloads. However, this hypothesis can and should be subjected to experimental verification.

Thus, by 2000 we expect that portions of this scenario will have occurred, that certain aircraft in certain airspace under certain conditions will be able to fly EFR. We would also expect that some CDTI-based EFR techniques will infiltrate more traditional IFR operations before then, assuming that CDTI systems become fairly common. However, we are not at all optimistic about placing control primarily in the cockpit by 2000.

We have dismissed this scenario for several reasons:

- It does not track well with the historical direction of ATC evolution. Ground-based control has been the norm for decades now. Stress upon this system is not severe enough to cause the radical departure implied by this scenario. Given the almost absolute requirement for evolutionary, gradual change within the system, revolutionary concepts stand little chance of entering the mainstream of ATC development.
- It is based upon highly uncertain technology. Automatic ATC problem-solving is still rudimentary, even in centralized applications. Introducing distribution merely complicates an already problematic situation.
- It is unlikely to be acceptable by either controllers or pilots. Controllers have already demonstrated their attitudes toward CDTI-based self-separation concepts. When asked to evaluate the safety of such operations in recent NASA-Ames laboratory work, they consistently rated EFR operations less safe than equivalent operations under their full control [11].
- Pilots have indicated that while they would almost uniformly approve of more and better cockpit traffic-avoidance aids, they also appreciate the safety advantages of having someone on the ground looking out for them. Doing away with routine surveillance ATC would leave matters solely to the pilots and automated aircraft, where failures that are considered rather minor today (e.g., avionics failures) could be catastrophic.

Although these primary deficiencies effectively disqualify this scenario from further consideration, EFR-based ATC does possess some interesting advantages:

- It could fill gaps in current ATC surveillance-based operations.

For instance, over water or mountainous terrain, where radar coverage is either nonexistent or spotty at best, EFR techniques could enable closer separations than are currently available using manual control techniques.

- Even highly limited implementations could greatly expand freedom of movement in the skies. Most en route flying is performed today under very tight constraints. Altitudes, routes, headings, and sometimes even speeds are assigned by controllers. Yet most en route time is spent nowhere near another aircraft. EFR techniques could permit pilots complete freedom of movement in those areas where current control techniques are unnecessary. Terminal control areas, of course, would continue to need the locally centralized, ground-based approach of today, as would highly congested en route airspace such as that over the northeastern United States.
- It may not be overly expensive. R&D costs will be considerable, of course, but if EFR operations were to replace the entire en route ground-based ATC system, equipment maintenance and personnel costs would be reduced significantly. If we assume that some form of CAS and CDTI systems will be constructed anyway during this time frame, the incremental cost of the R&D necessary to unify them under the control of an "intelligent-executive" cockpit black box may be quite small.

However, neither these advantages nor those stemming from the previous high-technology ATC scenario are sufficient to overcome the momentum now building for evolution of the current ground-based system.

Appendix C

EXPERIMENTS TO REDUCE THE UNCERTAINTIES

This appendix outlines a wide-ranging experimental program designed to resolve some of the uncertainties discussed in this report. The program is organized into phases ordered by time and complexity and requiring simulation facilities of varying degrees of sophistication. Each phase culminates in a usable interim system. This progression marries the already-begun AERA development process, with its attendant schedule of module development, with an incremental deployment regime. Thus, the experimental program is essentially scenario-independent, and it enables the earliest possible deployment of each module.

PHASE I

The initial experimental focus should be on the utility of a baseline controller planning aid. This system should include at least the following capabilities:

- Fuel-efficient initial planning of routes and delays.
- Plan alteration using altitude, course, and speed control.
- Graphic planning display showing static path trajectories.
- Textual display showing flight strips and clearance prompts.
- Digital air/ground communications.

To investigate the advantages and disadvantages of such an aid, the following experimental and analytic efforts are necessary:

1. Comparison of the effectiveness of individual, sequential planning against more complex forms of planning. Effectiveness, in this context, refers to such criteria as separation, fuel efficiency, schedule delays, and number of commands. The comparison should be made using high-fidelity testbed data under a variety of environmental conditions—airspace density, traffic mix, and weather. The already operational planning programs in the FAA testbed can be applied in real time to actual traffic tapes, resulting in performance statistics for individual planning. The more complex multi-aircraft plan-

ning algorithms can then be applied in non-real time to the same data to produce comparable performance estimates.

2. Determination of the human interface requirements at this level of development. Interface support for this interim system is expected to include displays of the airspace configuration, graphic planning displays for the path trajectories of the subject and object aircraft, and tabular displays alerting the controller to system errors or to requests for intervention. The main question here concerns the use of static vs. dynamic graphic planning displays. The abstracted control/display console should be used to compare time-stepped and event-stepped dynamic planning displays with the existing static planning display.
3. Determination of the human monitoring and control demands associated with each of the above planning algorithms and graphic planning displays. Task analysis studies are needed to determine the time lags and errors resulting from a take-over of control by the monitoring human. This can be accomplished using abstracted simulations and the FAA testbed.

One of the most pressing questions we must confront in the Shared Control scenario is embodied in item 3, i.e., the question of productivity gains when the human is still in the loop. We have therefore outlined an experiment designed to estimate productivity gains of several different Shared Control configurations.

In brief, the experiment compares two Shared Control configurations with a corresponding Baseline (non-automated) system. Productivity, as measured by the number of aircraft handled at a specified level of safety, will be determined under several different environmental conditions. To maximize validity, the experiments will use experienced controller teams and modified versions of the high-fidelity MITRE testbed.

The configurations we recommend for comparison are:

1. Baseline System—the current non-automated system under NAS Stage A procedures, upgraded to include DABS and ETABS capabilities.
2. Tactical-Only Shared Control System—the Baseline system plus automated capabilities for tactical conflict monitoring, tactical command generation, and digital clearance delivery.
3. Full Shared Control System—a more complete Shared Control system consisting of the above functions plus capabilities for automated profile generation and strategic planning.

We have chosen these three configurations because they include

the major aiding functions provided to the controller, and because they are relatively straightforward to implement and test. Other automated functions, such as flow control, resectorization, and failure recovery, contribute only marginally to the minute-by-minute operations of ATC. These added functions would require extensive development efforts, while most of the automated monitoring and planning capabilities listed in systems 2 and 3 have already been implemented in the MITRE testbed. Major changes are needed with respect to the human interface, though, including the addition of interfaces for controlling the graphic planning display, requesting system status data, and inputting parameter changes.

The proposed experiment is a two-stage process. The first stage involves implementing the Baseline and Tactical-Only Shared Control configurations and comparing productivities of the two under normal en route conditions. Experienced controllers will be assigned to two-person teams and trained in both the Baseline and Shared Control modes. Using a repeated-measures design, they will then experience both control conditions. It is expected that the two configurations will differ in both productivity and level of safety. In order to concentrate the variance into the productivity measure, the controller teams should be subjected to increasingly higher traffic loads until a set threshold of safety is violated. The final traffic loads (verified by testing above and below the final point) should be reliable indicators of productivity.

The second stage involves implementing the full Shared Control configuration and comparing its performance with that of the corresponding Baseline system. Here substantially larger productivity increases are expected. An experiment similar to that in the first stage should be performed, with several extenuating conditions. In the interest of determining the range of problems to which the system is applicable, additional conditions of adverse weather, restricted areas, and transition sector control should be included. Also, by changing the operator role from that of active participant to that of monitor with veto power over the automated functions, we can make an initial comparison of the productivity gains expected with the fully automated system and the Shared Control system.

PHASE II

The second experimental phase concerns automated out-of-track monitoring and more extensive replanning. It also injects more automated flow-control capabilities than mere delay computations. The following experimental tasks will be required:

1. Comparison of methods for automated out-of-track monitoring. The AERA system may continuously monitor tracks for out-of-threshold behavior or may establish critical points for checking, i.e., points at which an action is necessary to assure separation or otherwise guarantee safety. This comparison should be possible using analytic models of aircraft behavior. Alternatively, it can be made using the current MITRE testbed, once the association checking programs are in place.
2. Determination of the effectiveness of centralized metering and flow control (one center coordinates all sectors) against that of distributed flow control (each sector communicates anticipated outgoing loads to its neighbors). This may require extensive experimentation using multiple simulated sectors.
3. Testing of alternative forms of flow displays. The displays may show regional densities, flexibility of pathways, remaining routes, densities over time, projected hot spots, load factors as a percentage of maximum, etc. They may be graphic or textual, static or dynamic, just as for individual aircraft planning.
4. Failure analyses for the above modules, both individually and common-mode. This will result in specifications for system redundancy, inputs regarding displays of system status, and indications of necessary procedures for human backup. Much of this may be done using analytic models.
5. Stability analysis of the now multi-level system. Stability tests require impulse response measurements of the following inputs: new influxes of aircraft, changes to aircraft plans, airport closings, and changes in planning strategies. This should result in specifications of the system response time, resonances, and damping.
6. Experiments on maintenance of controller proficiency to determine necessary staffing levels and display requirements for failure backup. This involves checks on the controller's ability to assess current and predicted demands, system capacity, and the viability of alternate routes and backup plans. It may require extensive experimentation using multiple simulated control/display consoles.

The major goal of this set of experiments is to determine what portions of the controller's tasks can be effectively and reliably automated.

PHASE III

The final phase of the experiment builds on the above work by investigating automated sectorwide replanning and failure recovery capabilities. The major experimental efforts for this step are:

1. Testing and refining of sectorwide replanning programs. These programs operate in event of airport closure, adverse weather, emergency operations, etc. Because of the magnitude of this planning task, particularly efficient algorithms are necessary. Experimental tests need to be made of
 - Fast-time look-ahead with variable degrees of abstraction.
 - Automated recall of previous simulation results.
 - Pruning of low-confidence options.The complexity of this task and its degree of interaction with all other system operations demands that the full MITRE testbed system be used—multiple sectors, multiple forms of communication, all operational modules, etc. One of the major results of this experimental effort should be the definition of specific criteria of when to activate backup clearances and when to continue normal planning and replanning.
2. Development and testing of displays for informing the controller of altered sectorwide plans. We need to determine operator capability for offloading some of the planning from the automated system. This is dependent on the types of displays and controls provided. These tests are similar to the above in experimental requirements.
3. Tests of partial- and full-system failures. By individually implementing increased protection regions, emergency flow control, plan coasting, and resectorization we can determine their contributions to failure backup. This should require the full MITRE testbed. We also need to test alternative displays for showing AERA system status: degree of system stabilization, predicted down time, number of operational modules, and status of coasting clearances.

An experimental program of this type could answer many of the questions raised in the preceding sections—questions concerning the proper functions for automation, the role of the human controller, and the expected levels of system performance and reliability. It should also result in tested, useful interim products leading to a final highly automated system.

REFERENCES

1. Amato, R. A., and W. A. Manison, "Advancements in ATC Information Displays and Controller-Computer Interaction," *Proceedings of 22nd Annual Meeting of Air Traffic Control Association*, October 10-13, 1977, pp. 316-323.
2. Klass, P. J., "Decision Computers Studied for ATC," *Aviation Week and Space Technology*, November 19, 1979.
3. Safeer, H. B., "Nontechnological Alternatives for Balancing Airport/Airspace Supply and Demand," *The FAA Plans and Programs for the Future Airport and Air Traffic Control System*, Office of Technology Assessment Seminar, November 13, 1980, U.S. Department of Transportation, pp. 152-168.
4. "FAA Readies Emergency Plan in Event of Strike by PATCO," *Aviation Week and Space Technology*, October 6, 1980, pp. 34-35.
5. Bowers, A., and R. Rucker, *A More Fully Automated En Route ATC System*, The MITRE Corporation, METREK Division, 1976.
6. Rucker, R., *Automated En Route ATC (AERA): Briefing Charts on Proposed Goals and R&D Approach*, The MITRE Corporation, Working Paper, May 1980.
7. Goldmuntz, L., et al., *Outline for AERA Concept Document*, FAA Office of Systems Engineering Management, December 1980.
8. Bowers, A. W., *AERA: A Proposed Man-Machine Interface to Support a Postulated New Role for the En Route Air Traffic Controller*, The MITRE Corporation, WP-79W663, October 1979.
9. Mertes, F., and L. Jenney, *Automation Applications in an Advanced Air Traffic Management System, Volume III: Methodology for Man-Machine Task Allocation*, prepared for the Department of Transportation by TRW Incorporated, Final Report No. DOT-TSC-OST-74-14, August 1974.
10. Kreifeldt, J., L. Parkin, P. Rothschild, and T. Wempe, "Implications of a Mixture of Aircraft With and Without Traffic Situation Displays for Air Traffic Management," presented at the Twelfth Annual Conference on Manual Control, University of Illinois, Urbana, Illinois, and Ames Research Center, Moffett Field, California, May 1976.

11. Del Balzo, J., "ATC in the Year 2000," presented at Air Traffic Control Association's 22nd Annual Meeting and Technical Program, Las Vegas, Nevada, October 10-13, 1977, *Supplement to Proceedings*, pp. 41-63.
12. Couluris, G., M. Tashker, and M. Penick, *Policy Impacts of ATC Automation: Human Factors Considerations*, prepared for the Department of Transportation by SRI International and Payne-Maxie Consultants, Report No. FAA-AVP-78-1, January 1978.
13. U.S. Department of Transportation, *Advanced Air Traffic Management System Study Overview*, DOT-TSC-75-32, June 1975.
14. Poritzky, S., "ATC Automation—The Dilemma of Automated Decisionmaking," presented at Air Traffic Control Association, Atlantic City, New Jersey, October 1979.
15. Whitfield, D., R. Ball, and G. Ord, "Some Human Factors Aspects of Computer-Aiding Concepts for Air Traffic Controllers," *Human Factors*, Vol. 22, No. 5, 1980, pp. 569-580.
16. Poritzky, S., "The Roadmap to the System of the Future," *The FAA Plans and Programs for the Future Airport and Air Traffic Control System*, U.S. Department of Transportation, Office of Technology Assessment Seminar, November 13, 1980, pp. 94-105.
17. Thackray, R. I., J. Bailey, and R. M. Touchstone, "Physiological, Subjective, and Performance Correlates of Reported Boredom and Monotony While Performing a Simulated Radar Control Task," in R. R. Mackie (ed.), *Vigilance, Theory, Operational Performance and Physiological Correlates*, Plenum Press, New York and London, 1977.
18. *Human Factors: Air Traffic Control I and II*, Vol. 22, Special Issues, Nos. 5 and 6, October and December 1980.
19. Kinney, G., M. Spahn, and R. Amato, "The Human Element in Air Traffic Control: Observations and Analyses of the Performance of Controllers and Supervisors in Providing ATC Separation Services," The MITRE Corporation, Technical Report MTR-7655, December 1977.
20. Koenke, E. J., "Long Term ATC System Improvements," *The FAA Plans and Programs for the Future Airport and Air Traffic Control System*, U.S. Department of Transportation, Office of Technology Assessment Seminar, November 13, 1980, pp. 140-151.
21. Andrews, J. W., and W. M. Hollister, "Electronic Flight Rules: An Alternative Separation Assurance Concept," Lincoln Laboratory Project Report FAA-RD-80-2, January 1980.

22. Steeb, R., S. Cammarata, F. Hayes-Roth, and R. Wesson, *Distributed Intelligence for Air Fleet Control*, The Rand Corporation, R-2728-ARPA, August 1981.
23. Wesson, R. B., *Planning in the World of the Air Traffic Controller*, Doctoral Dissertation, Department of Computer Sciences, University of Texas at Austin, 1977.

END
DATE
FILMED
4-82

DTIC